Peaks vs. Components

Sebastian Vollmer[*]
Hajo Holzmann[†]
Florian Schwaiger[‡]

Leibniz Universität Hannover, Discussion Paper No. 452
ISSN 0949-9962

We analyze the cross-national distribution of GDP per capita and its evolution from 1970 to 2009. We argue that peaks are not a suitable measure for distinct convergence clubs/equilibria in the cross-country distribution of GDP per capita, because the number of peaks is not invariant under non-linear strictly monotonic transformations of the data such as the logarithmic transformation. Instead, we model the distribution as a finite mixture, and determine its number of components via statistical testing. We find that the number of components in the cross-country distribution changes from three to two in the mid 1990s.

JEL classification: C12, O11, O47, F01
Keywords: twin peaks, economic growth, convergence.

[*]University of Göttingen & Harvard University
[†]University of Marburg
[‡]University of Marburg

1 Introduction

The notion of twin peaks in the cross-country distribution of GDP per capita was introduced by Quah (1993, 1996, 1997). He interpreted the emergence of twin peaks as polarization distribution into a rich and a poor convergence club. Bianchi (1997) confirmed Quah's observation of twin peaks via rigorous statistical testing. The contributions of Quah are part of a larger literature on convergence (e.g. Barro, 1991; Barro and Sala-i-Martin, 1992; Mankiw, Romer and Weil, 1992; Sala-i-Martin, 1996; Galor, 1996; Jones, 1997; Graham and Temple, 2006). It is controversial in this literature whether the twin peaks represent locally stable equilibria/convergence clubs (Quah, 1996) or whether they are only a temporary phenomenon due to a high frequency of growth miracles (Jones, 1997).

The unified growth theory (c.f. Galor, 2010 for an overview) provides another explanation for multiple regimes in the cross-country distribution of GDP per capita which also uncovers the forces that have lead to the emergence of these regimes. The theory suggests that growth segments economies into three fundamental regimes: a malthusian regime with slow growing economies, fast growing economies in a sustained growth regime, and a third group in the transition from one regime to the other. One important difference to models with multiple equilibria is that this segmentation does not represent the long-run steady state of these economies. Variations in the levels of income only reflect country-specific characteristics and not the actual stage of development. Thus, there are no critical levels that permit economies to switch from one regime to the other, but rather critical rates of progress.

Recent theoretical work by Schumacher (2009) and Strulik (2012) provides alternative explanations for the emergence of multiple equilibria in the cross-country distribution of GDP per capita. Schumacher (2009) endogenizes discounting via wealth in a neoclassical growth model and shows that this can generate multiple equlibria. Strulik (2012) formulates an endogenous growth theory with endogenous patience, which can explain the take-off from stagnation to modern growth. He concludes that either all countries adjust to the same balanced growth path or that lagging countries will never catch up.

In this paper we challenge Quah's twin peaks result. We show that the number of peaks of a distribution is not preserved under strictly monotonic transformations of the data: a simple log transformation may change the number of peaks in the cross-country distribution of GDP per capita. This fact casts doubts on the economic interpretation of twin peaks: it does not make much sense to call a country high income on the original scale of the GDP per capita data and middle income on a log scale of the same data. A suitable measure of convergence clubs or growth regimes should not be affected by a simple log transformation of the data.

We therefore propose another method to identify different regimes within a distribution which does not have this problem. We will use mixture models to estimate the cross-country distribution of GDP per capita and to statistically assign countries to different convergence clubs. Mixture models are not new to the economic literature and have been used in quite a few articles to model income distributions. Most prominently, Paap and Dijk (1998) used a two-component mixture, consisting of a truncated normal distribution and a Weibull distribution, to model the cross-country distribution of GDP per capita. Nevertheless, we are not aware of any article that challenges the twin peaks approach and suggests mixture models as an alternative.

2 Data

We use the Penn World Tables 7.0 (PWT) data for the period from 1970 to 2009. The PWT is a panel dataset containing 190 countries and 38 variables. We use the variable rgdpch, which is PPP converted GDP per capita (chain series) at 2005 constant prices. We consider the mentioned GDP per capita variable on its original scale ($1,000) and on a logarithmic scale with base 10.

We exclude a few small countries whose economies heavily depend on oil export: Bahrain, Brunei, Equatorial Guinea, Gabon, Kuwait, Qatar, Suriname, Timor-Leste and Trinidad and Tobago from the analysis. The reason for this choice is, that these countries show large fluctuations in GDP per capita, which are mostly driven by fluctuations of the oil price. Arguably, these countries are not relevant for understanding multiple equilibria in the world's cross-country distribution of GDP per capita . The PWT dataset contains two versions of China, we thus exclude the second version (CH2) from the analysis. We believe that using a balanced panel is the most appropriate to analyze the cross-country distribution of GDP per capita over time, because a balanced panel is not affected by changes in the sample composition. This leaves 151 countries in the dataset, for which we have GDP per capita data for all years from 1970 to 2009.

3 Peaks

Figure 1 shows simple kernel density estimates of the cross-country distribution of GDP per capita in 1985 on the original scale ($1000) and on a logarithmic scale with base 10. The density of the data on the original scale has two peaks and the density of the data on a log scale has three peaks. This simple picture illustrates that the number of peaks is not preserved under a simple log transformation: Quah's twin peaks become triple peaks on the log scale.

However, the different numbers of peaks in the plots could be a simple artifact of the nonparametric curve estimates, e.g. from inaccurate choice of the tuning parameter. It is therefore necessary to validate the statistical significance of the peaks via rigorous statistical testing. To this end we utilize Silverman's test. Formally, a *peak* of a density f (and similarly of the kernel estimator $\hat{f}$) is a local maximum of f (or $\hat{f}$). Silverman (1981) showed that the number of modes of $\hat{f}$ is a right-continuous, monotonically decreasing function of the bandwidth h if the normal kernel $K(x) = (2\pi)^{-1}\exp(-x^2/2)$ is employed. This allowed him to define the k-critical bandwidth $h_c(k)$ as the minimal bandwidth h for which $\hat{f}$ still just has k modes and not yet $k+1$ modes. Based on the notion of the

k-critical bandwidth, Silverman (1981) proposed a bootstrap test for the hypotheses

$$\tilde{H}_k : f \text{ has at most } k \text{ modes} \qquad \text{against} \qquad \tilde{K}_k : f \text{ has more than } k \text{ modes.}$$

This test is known to be slightly conservative (even asymptotically), for $\tilde{H}_1$ we therefore use the adjustment proposed by Hall and York (2001). The tests were performed using our R-package `silvermantest`[1] We apply Silverman's test to the distributions of GDP per capita and log-GDP per capita for all years from 1970 to 2009. We report the p-values in Tables 1 and 2. The left hand column of the tables shows the different null hypotheses of the tests.

For the distribution of GDP per capita we can reject the null hypothesis of a single peak from 1970 to 1990, but we cannot reject the null hypothesis of two peaks in favor of three or more peaks. This is basically the period that Quah studied in his influential papers, and our results confirm his findings. From 1991 onwards we can also reject the null hypothesis of two peaks in favor of three peaks (but not more). Thus, we find evidence for two peaks from 1970 to 1990 and for three peaks thereafter. For the distribution of log-GDP per capita we can reject the null hypothesis of two peaks in favor of three peaks (but not more) from 1970 to 1990, but we fail to reject the null hypothesis of a single peak. Note that this result does not mean that the null hypothesis of a single peak is correct, it just means that there is not enough evidence to reject it at a level of 5%. Thus, there is evidence of three peaks, but none of only two peaks from 1970 to 1990. From 1991 onwards we cannot reject any of the null hypotheses[2] and thus find evidence for only a single peak.

What do we learn from this analysis? The number of peaks is relevant information for the proper visualization of data. However, our results show clearly that peaks should neither be used for economic interpretation of the cross-country distribution of GDP per capita nor for assigning countries to convergence clubs, growth regimes and the like. It does not make sense to conclude that the distribution of the GDP per capita consists of two convergence clubs between 1970 and 1990, while the distribution of log GDP per capita consists of three convergence clubs over the same period.

4 Components

4.1 Methods

We now turn to mixture models to estimate the cross-country distribution of GDP per capita. Let f_X denote the density of the cross-country distribution of GDP per capita X for a given year. We model

$$f_X(x) = \alpha_1 g(x; \phi_1) + \ldots + \alpha_m g(x, \phi_m), \qquad x > 0,$$

[1]available online at http://www.uni-marburg.de/fb12/stoch/research/rpackage

[2]with the exception of a few transition years from 1992 to 1994 where the distribution appears to have four peaks

where $g(x;\phi)$ is a parametric family of densities and the weights $\alpha_i \geq 0$ sum up to one. There is no general simple connection between the number of modes of f and the number of components m. Typically, for unimodal g, the number of modes of f will be at most m, but often will be less than m. The number of components is preserved if the data are transformed via a strictly monotonic transformation (if densities are correspondingly transformed). We let $Y = \log X$ and model the density of log-income f_Y by

$$f_Y(y) = \alpha_1\varphi(y;\mu_1,\sigma_1) + \ldots + \alpha_m\varphi(y;\mu_m,\sigma_m),$$

where $\varphi(\cdot,\mu,\sigma)$ is the density of the normal distribution with mean $\mu \in \mathbb{R}$ and standard deviation $\sigma > 0$. Then $g(\cdot;\phi)$ in the representation of f_X is the log-normal distribution. The number of components is determined via statistical inference: Ee aim to test successively for ascending m in $\mathbb{N}$ the hypotheses

$$H_{m_0} : m_0 = m \quad \text{against} \quad K_{m_0} : m_0 \geq m+1,$$

where $m_0 \in \mathbb{N}$ is the true, unknown number of components. Testing in parametric models is often accomplished by using the likelihood ratio test (LRT). However, the standard theory of the LRT does not apply for the number of components in finite mixture models (Dacunha-Castelle and Gassiat, 1999). Recently, Chen et al. (2001, 2004) and Chen and Kalbfleisch (2005) suggested modified LRTs, which retain comparatively simple limit theory as well as the good power properties of the LRT. Unfortunately, these tests are only valid if the switching parameter is one-dimensional and hence we cannot apply them for selecting the number of components.

In our setting with switching μ and σ only an asymptotic test for homogeneity, i.e. for H_1, is available, see Chen and Li (2009). Therefore, in order to test all hypotheses under investigation with the same methodology, we apply the commonly used parametric bootstrap.[3] As is well known, since μ and σ both switch the likelihood function is unbounded if small values of the standard deviation are allowed. Therefore, we use a penalized log-likelihood as proposed in Chen and Li (2009) as follows:

$$l_n(X_1,\ldots,X_n;\boldsymbol{\mu},\boldsymbol{\sigma},\boldsymbol{\alpha}) = \sum_{i=1}^{n} \log\Big(\sum_{j=1}^{m} \alpha_j\varphi(X_i;\mu_j,\sigma_j)\Big) + \cdot p_n(X_1,\ldots,X_n,\boldsymbol{\sigma}),$$

where $\boldsymbol{\mu} = (\mu_1,\ldots,\mu_m)$, $\boldsymbol{\sigma} = (\sigma_1,\ldots,\sigma_m)$ and $\boldsymbol{\alpha} = (\alpha_1,\ldots,\alpha_{m-1})$ with $\alpha_m := 1 - \sum_{j=1}^{m-1}\alpha_j$ for $m > 1$ and $\boldsymbol{\alpha} = 1$ for $m = 1$, and

$$p_n(X_1,\ldots,X_n,\boldsymbol{\sigma}) = -\frac{1}{20}\sum_{j=1}^{m}\big(s_n^2/\sigma_j^2 + \log(\sigma_j^2/s_n^2)\big),$$

where $s_n^2 = n^{-1}\sum_{i=1}^{n}(X_i - \bar{X})^2$ with $\bar{X} = n^{-1}\sum_{i=1}^{n}X_i$. The function p_n penalizes small values of a σ_j, and guarantees a bounded (penalized) likelihood.

[3] We used 1000 bootstrap replications.

After fitting the model and selecting the number of components, we can use the mixture model for cluster analysis, see e.g. Fraley and Raftery (2002). Each observation can be assigned posterior probabilities to belong to each of the components in the mixture model. This yields three levels of income which we label poor, intermediate and rich, with indices $1, 2, 3$ when a three component mixture is fitted or two levels of income which we label poor and rich when a two component mixture is fitted. Specifically, the posterior probability of an observation y to belong to group j is equal to

$$p_j(y) = \frac{\hat{\alpha}_j \varphi(y; \hat{\mu}_j, \hat{\sigma}_j)}{\hat{\alpha}_1 \varphi(y; \hat{\mu}_1, \hat{\sigma}_1) + \ldots + \hat{\alpha}_m \varphi(y; \hat{\mu}_m, \hat{\sigma}_m)},$$

for $m = 2$ or $m = 3$ in case of a two- or a three-component mixture. Therefore, we do not merely assign an income level to each country, but rather a probability distribution, which makes transitions from one group to the other much more transparent. One may then assign an observation y to one of the components by using the maximum a-posterior estimate (MPE), which assigns the j to the country i for which $p_j(y)$ is maximal. One can also determine the threshold $t_{j,j+1}$, $j = 1, \ldots, m-1$, for the values of the log-GDP per capita at which the MPE changes between group j and $j+1$, by solving the equations

$$p_j(t_{j,j+1}) = p_{j+1}(t_{j,j+1}),$$

restricted to the interval $[\hat{\mu}_j, \hat{\mu}_{j+1}]$.

4.2 Results

Table 3 displays the results of the parametric bootstrap test based on 1000 bootstrap samples. We can always reject the hypothesis of homogeneity, i.e. of a single normal distribution. Further, we cannot reject the null hypothesis of two components in 1970, 1971 and 1972 at the 5 percent significance level, however, the p-values are already quite. From 1973 to 1995 we can reject the null hypothesis of two components with p-values at the 5% level. From 1996 to 2001 the p-values are still quite low, but we cannot reject the null hypothesis at the 5 % level anymore. After 2002 the p-values are rather large and the null hypothesis cannot be rejected. Overall, we observe a three component mixture that evolves into a two component mixture. We thus model the cross country distribution of GDP per capita with three components from 1970 to 1995 and with two components from 1996 to 2009.

In Figure 2 we show the fitted three-component mixtures for 1975 and 1985 and compare it to the corresponding kernel density estimators based on the smallest bandwidths which produce three modes. Further, Figure 3 shows the fitted two-component mixtures for 1996 and 2005 with the corresponding kernel density estimators based on the smallest bandwidths which produce two modes. We also provide quantile-quantile (qq) plots of the data against the fitted mixture models, see figure 4 and 5. The qq-plots show that the three respectively the two component mixtures describe the data well.

Figure 6 shows the development of the different component means over time as well as the thresholds where the maximum a-posterior estimate changes from one component to the other. The component means are also shown in Table 4. The mean of the low-income and the middle-income component hardly changes between 1970 and 1990, but both component means show substantial increases from 1991 to 1995. The mean of the high-income component steadily grows from 1970 to 1995 (by roughly 50 percent over the entire period).

After 1995 the three components merge into two components. The new higher-income component basically continues on the growth path of the high-income component from the previous model, whereas, the low-income and middle-income components from the previous model merge into a new lower-income component. Both component means steadily grow between 1996 and 2009 (both roughly by one third over the entire period).

The observation that the low-income and middle-income components of the three-component model merge into a new lower-income component in the two-component model is also supported by the component weights which are displayed in Table 5). In 1970 the low-income component constitutes about 50 percent of the countries, whereas the middle-income and high-income components represent 33 and 17 percent respectively. Over time, this picture reverses: Between 1970 and 1990, the size of the low-income component decreases to roughly 31 % and the size of the middle income component increases to 50 %. After 1996 the lower-income component is about as large as the low-income and middle-income components were jointly, which again supports the observation that those two components merged into a new lower-income component. Between 1991 and 1995 there is some variation in the component sizes, but before and after this picture is remarkably stable.

It is also important to keep the relative component sizes in mind when we interpret the component means. Even though the means of both the low-income and middle-income components stagnated between 1970 and 1990, there was still quite a bit of growth, because many countries made transition from the low-income component to the middle income component.

5 Concluding Remarks

In this paper we challenge the long standing twin peaks finding in the cross-country distribution of GPD per capita. We show that the number of peaks of a distribution depends on the scale (e.g. original or logarithmic) and argue that this feature is highly undesirable for economic interpretations. As an alternative approach to peaks, we use finite mixture models to investigate the cross-country distribution of GDP per capita, since *a.* their number of components does not depend on the scale, *b.* components in the mixture arguably correspond better to income clubs in the distribution than peaks, and *c.* finite mixture models allow for an accurate analysis of the intra-distributional dynamics by using posterior probability estimates.

Interestingly, our conclusions are not so different from Quah's, however, this might well be a coincidence. For the period that Quah studied, we find that the cross-country distribution of GDP per capita consisted of three components, which seem more like transition regimes rather than convergence clubs. Only for more recent years we find that the cross-country distribution of GDP per capita consists of two groups which are quite stable and follow their own growth paths, and thus could potentially be interpreted as convergence clubs. In any case, we wanted to make the point that in our opinion, components should take the place of peaks in the literature on economic growth, because they do not suffer from the inherent shortcomings that peaks have and thus can lead to more meaningful economic interpretations.

References

Barro, R. (1991) Economic Growth in a Cross Section of Countries. *Quarterly Journal of Economics* **106**: 407–443.

Barro, R. and X. Sala-i-Martin (1992) Convergence. *Journal of Political Economy* **100**: 223–251.

Bianchi, M. (1997) Testing for Convergence: Evidence from Non-Parametric Multimodality Tests. *Journal of Applied Econometrics* **12**: 393–409.

Chen, H., Chen, J. and Kalbfleisch, J. D. (2001) A modified likelihood ratio test for homogeneity in finite mixture models. *Journal of the Royal Statistical Society Series B* **63**: 19–29.

Chen, H., Chen, J. and Kalbfleisch, J. D. (2004) Testing for a finite mixture model with two components. *Journal of the Royal Statistical Society Series B* **66**: 95–115.

Chen, J. and Kalbfleisch,J. D. (2005) Modified likelihood ratio test in finite mixture models with a structural parameter. *Journal of Statistical Planning and Inference* **129**: 93–107.

Chen, J. and Li, P. (2009) Hypothesis test for normal mixture models: the EM approach. *Annals of Statistics* **37**: 2523–2542.

Dacunha-Castelle, D. and Gassiat, E. (1999) Testing the order of a model using locally conic parametrization: population mixtures and stationary ARMA processes. *Annals of Statistics* **27**: 1178–1209

Fraley, C. and Raftery, A. E. (2002) Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association* **97**: 611–631.

Galor, O. (1996) Convergence? Inference from Theoretical Models. *Economic Journal* **106**: 1056–1069.

Galor, O. (2010) The 2008 Klein Lecture: Comparative Economic Development: Insights from Unified Growth Theory. *International Economic Review* **51**: 1–44.

Graham, B. S. and Temple, J. R. W. (2006) Rich nations, poor nations: how much can multiple equilibria explain? *Journal of Economic Growth* **11**: 5–41.

Hall, P. and York, M. (2001) On the calibration of Silverman's test for multimodality. *Statistica Sinica* **11**: 515–536.

Heston, A.; Summers, R. and Aten, B. (2011) Penn World Tables Version 7.0. *Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.*

Jones, C. (1997) On the Evolution of the World Income Distribution. *Journal of Economic Perspectives* **11**: 19–36.

Mankiw, G., D. Romer and D. Weil (1992) A Contribution to the Empirics of Economic Growth. *Quarterly Journal of Economics* **107**: 407–437.

Paap, R. and Dieck, H. K. (1998) Distribution and mobility of wealth of nations. *European Economic Review* **42**: 1269–1293.

Quah, D. T. (1993) Galton's Fallacy and Tests of the Convergence Hypothesis. *Scandinavian Journal of Economics* **95**: 427–443.

Quah, D. T. (1996) Twin Peaks: Growth and Convergence in Models of Distribution Dynamics. *Economic Journal* **106**: 1045–1055.

Quah, D. T. (1997) Empirics for growth and distribution: Stratification, polarisation and convergence clubs. *Journal of Economic Growth* **2**: 27-59.

Sala-i-Martin, X. (1996) The Classic Approach to Convergence Analysis. *Economic Journal* **106**: 1019–1036.

Schumacher, I. (2009) Endogenous discounting via wealth, twin-peaks and the role of technology, *Economics Letters*, **103** : 78-80

Silverman, B. W. (1981) Using kernel density estimates to investigate multimodality. *Journal of the Royal Statistical Society Series B* **43**: 97–99.

Strulik H. (2012) Patience and Prosperity *Journal of Economic Theory* **147**: 336–352.

Tables and Figures

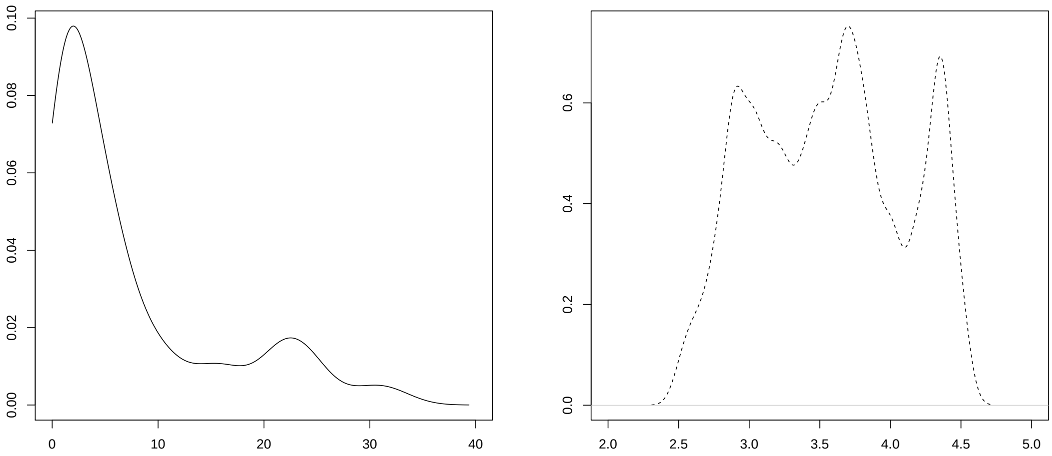

(a) GDP per capita in \$1000 (b) log GDP per capita

Figure 1: Kernel density estimate for GDP per capita in \$1000 (left) and log GDP per capita (right) for 1985. We use the logarithm to the base 10.

	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979
at most 1	0.000	0.000	0.000	0.000	0.000	0.000	0.001	0.030	0.023	0.008
at most 2	0.267	0.245	0.286	0.170	0.154	0.429	0.932	0.814	0.954	0.916
at most 3	0.585	0.644	0.620	0.768	0.874	0.812	0.858	0.528	0.740	0.854
at most 4	0.170	0.232	0.238	0.624	0.775	0.678	0.728	0.770	0.352	0.596
	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989
at most 1	0.001	0.002	0.001	0.000	0.000	0.000	0.007	0.000	0.001	0.000
at most 2	0.650	0.302	0.672	0.806	0.430	0.364	0.390	0.784	0.833	0.508
at most 3	0.569	0.251	0.814	0.598	0.316	0.038	0.042	0.476	0.361	0.523
at most 4	0.546	0.954	0.836	0.520	0.784	0.750	0.488	0.152	0.184	0.175
	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999
at most 1	0.003	0.011	0.021	0.044	0.041	0.016	0.009	0.011	0.021	0.061
at most 2	0.251	0.024	0.017	0.002	0.008	0.031	0.142	0.138	0.042	0.002
at most 3	0.434	0.460	0.038	0.022	0.009	0.065	0.178	0.247	0.407	0.296
at most 4	0.326	0.124	0.134	0.106	0.222	0.254	0.245	0.116	0.188	0.220
	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
at most 1	0.031	0.029	0.026	0.036	0.029	0.025	0.011	0.006	0.005	0.009
at most 2	0.000	0.000	0.002	0.000	0.000	0.002	0.008	0.017	0.042	0.034
at most 3	0.300	0.136	0.552	0.580	0.638	0.676	0.608	0.310	0.221	0.154
at most 4	0.221	0.101	0.536	0.235	0.278	0.378	0.416	0.144	0.022	0.204

Table 1: P-values for testing the number of peaks in the cross-country distribution of GDP per capita with Silverman's test.

	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979
at most 1	0.215	0.212	0.129	0.011	0.033	0.104	0.308	0.551	0.393	0.529
at most 2	0.084	0.051	0.003	0.006	0.013	0.024	0.034	0.130	0.075	0.118
at most 3	0.202	0.352	0.413	0.858	0.824	0.360	0.144	0.054	0.048	0.011
at most 4	0.729	0.596	0.656	0.654	0.367	0.290	0.711	0.276	0.822	0.850
	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989
at most 1	0.533	0.163	0.249	0.167	0.197	0.221	0.180	0.246	0.264	0.264
at most 2	0.128	0.048	0.039	0.034	0.018	0.011	0.013	0.020	0.036	0.012
at most 3	0.002	0.032	0.063	0.699	0.859	0.956	0.607	0.850	0.734	0.404
at most 4	0.811	0.822	0.414	0.770	0.780	0.816	0.744	0.642	0.764	0.348
	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999
at most 1	0.220	0.220	0.228	0.275	0.246	0.227	0.173	0.168	0.139	0.103
at most 2	0.008	0.166	0.118	0.180	0.128	0.352	0.386	0.235	0.202	0.154
at most 3	0.878	0.004	0.013	0.013	0.120	0.214	0.408	0.166	0.389	0.388
at most 4	0.708	0.840	0.457	0.605	0.199	0.070	0.174	0.848	0.594	0.091
	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
at most 1	0.136	0.105	0.102	0.109	0.133	0.126	0.157	0.171	0.245	0.262
at most 2	0.208	0.140	0.068	0.122	0.173	0.112	0.112	0.158	0.117	0.104
at most 3	0.048	0.326	0.259	0.475	0.544	0.667	0.566	0.492	0.563	0.722
at most 4	0.126	0.284	0.444	0.472	0.138	0.250	0.242	0.258	0.261	0.436

Table 2: P-values for testing the number of peaks in the cross-country distribution of log GDP per capita with Silverman's test.

	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979
1 vs. 2	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
2 vs. 3	0.07	0.16	0.16	0.01	0.01	0.02	0.01	0.02	0.07	0.02
	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989
1 vs. 2	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
2 vs. 3	0.02	0.01	0.00	0.01	0.02	0.01	0.01	0.02	0.04	0.02
	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999
1 vs. 2	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
2 vs. 3	0.01	0.04	0.41	0.07	0.07	0.04	0.06	0.19	0.06	0.05
	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009
1 vs. 2	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
2 vs. 3	0.16	0.10	0.05	0.27	0.40	0.23	0.33	0.38	0.41	0.29

Table 3: Bootstrap p-values for testing the hypotheses of one and two components in the cross-country distribution of GDP per capita.

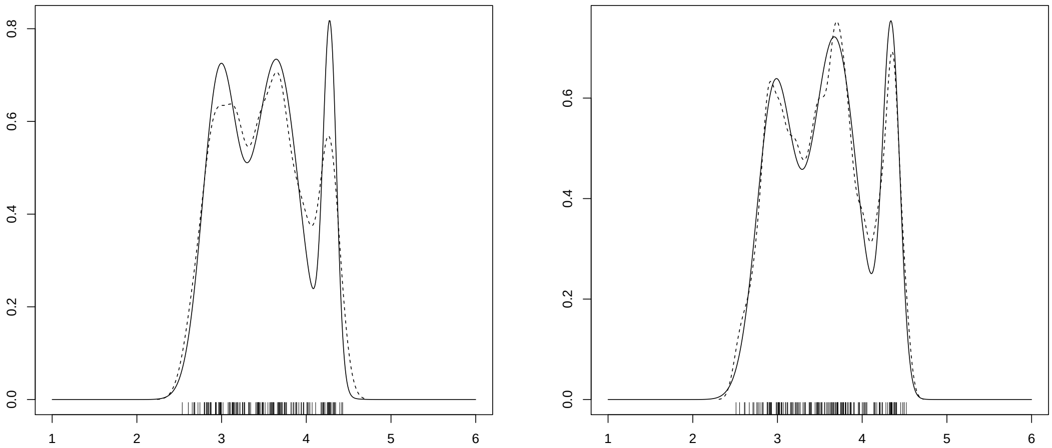

(a) 1975 (b) 1985

Figure 2: Fitted three-component mixture densities (solid line) and kernel density estimate based on $h_c(3)$ (dashed line) for the log-data.

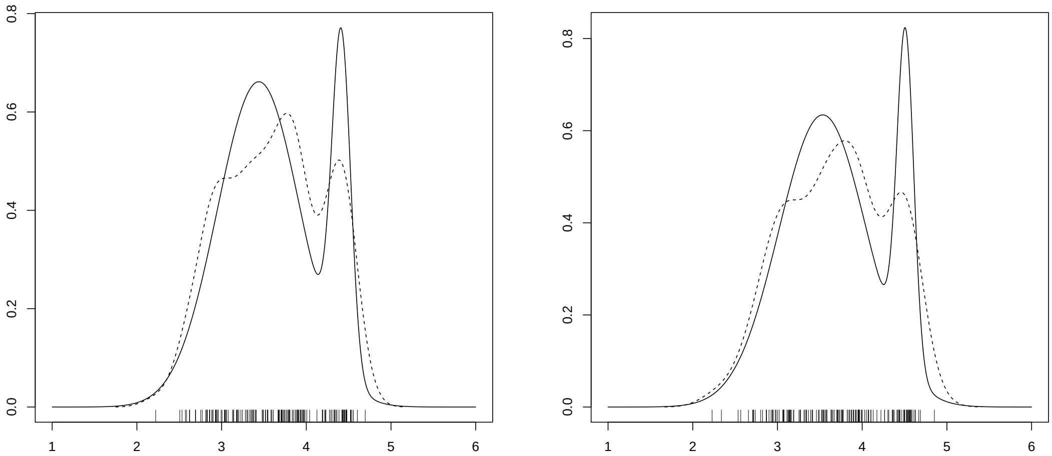

(a) 1996 (b) 2005

Figure 3: Fitted two-component mixture densities (solid line) and kernel density estimate based on $h_c(2)$ (dashed line) for the log-data.

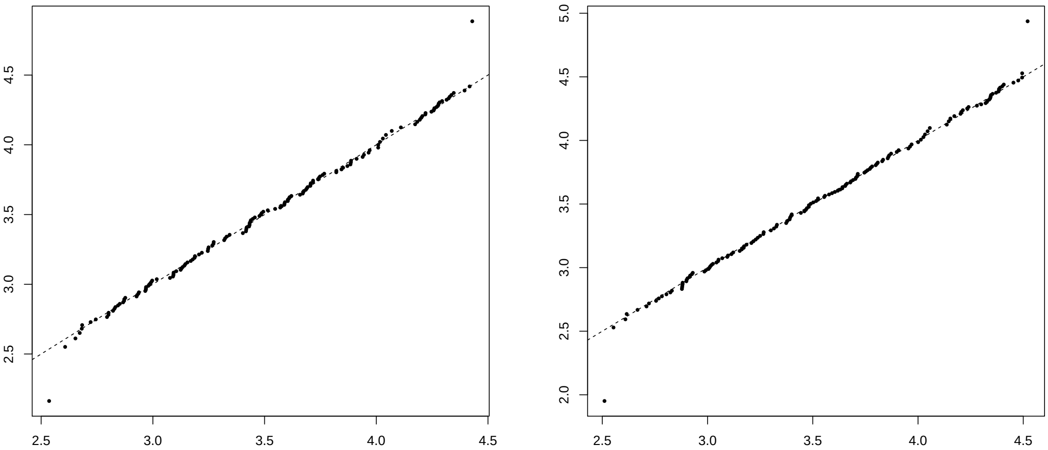

(a) 1975 (b) 1985

Figure 4: QQ plot of the log-data with three components.

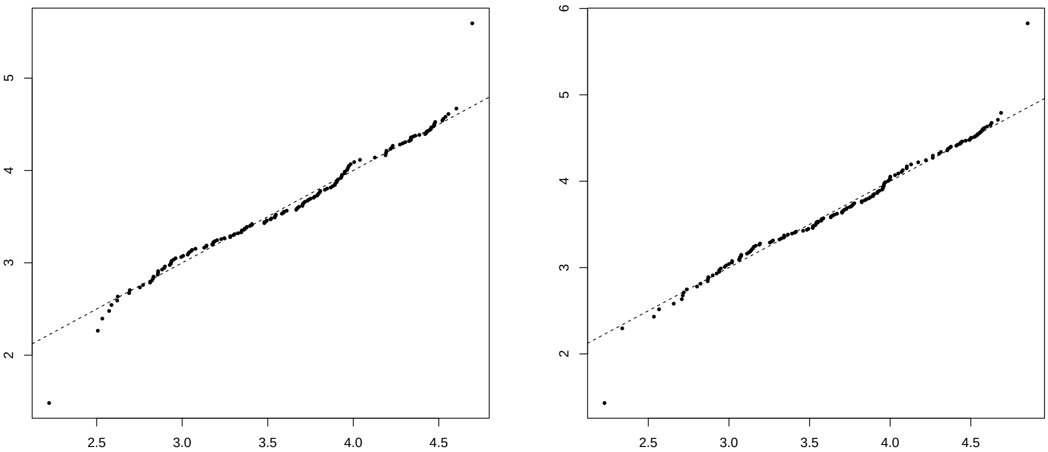

(a) 1996 (b) 2005

Figure 5: QQ plot of the log-data with two components.

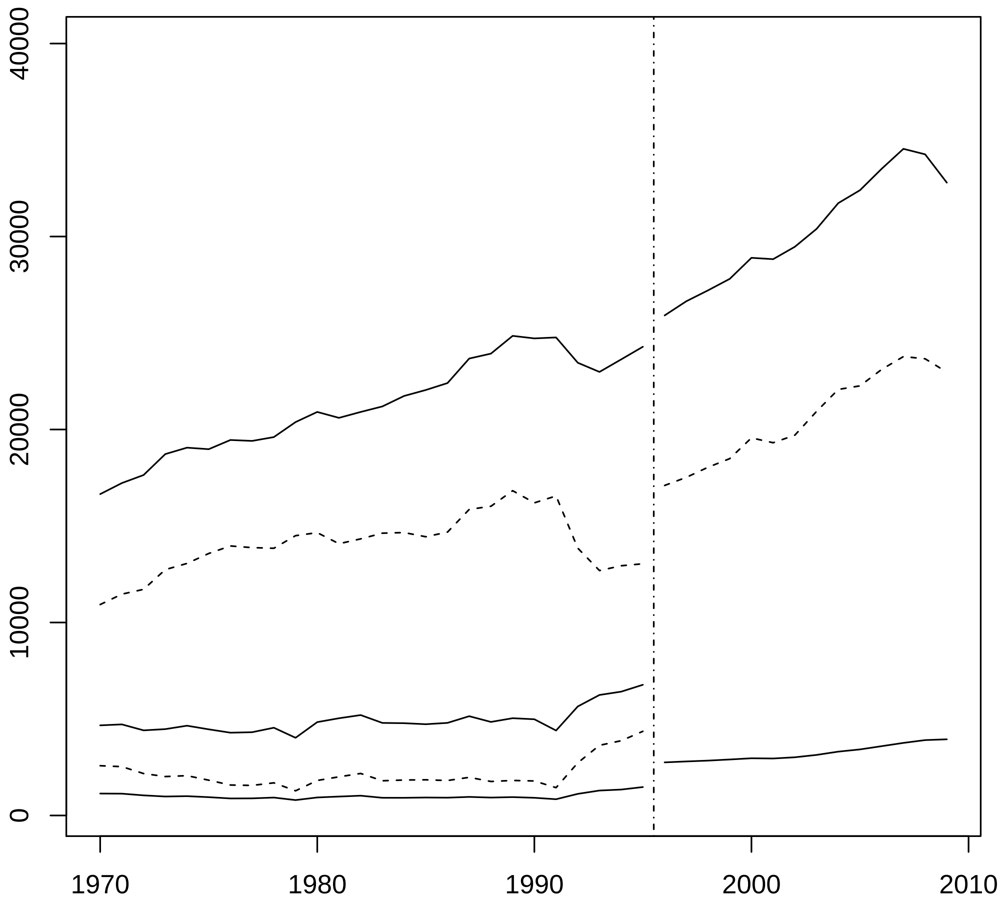

Figure 6: Estimated means (based on the log GDP data, but displayed on the original scale) of the three or respectively two distinct groups (solid lines). Income levels where the maximum a-posterior estimates switch from one group to the other (dashed lines).

	low	middle	high
1970	1136	4669	16651
1971	1129	4719	17225
1972	1042	4411	17641
1973	983	4475	18726
1974	1000	4652	19058
1975	948	4463	18980
1976	881	4290	19457
1977	886	4314	19408
1978	928	4547	19607
1979	798	4025	20385
1980	933	4838	20910
1981	979	5038	20604
1982	1025	5202	20908
1983	916	4796	21196
1984	915	4782	21739
1985	930	4728	22048
1986	922	4798	22407
1987	962	5144	23678
1988	927	4847	23933
1989	952	5040	24853
1990	916	4985	24721
1991	842	4401	24771
1992	1120	5649	23460
1993	1292	6245	22985
1994	1342	6416	23634
1995	1472	6773	24290

(a) Component means for the years 1970 to 1995 (balanced dataset).

	low	high
1996	2752	25910
1997	2799	26646
1998	2842	27211
1999	2901	27807
2000	2964	28897
2001	2953	28827
2002	3016	29470
2003	3137	30396
2004	3308	31732
2005	3423	32402
2006	3591	33512
2007	3760	34544
2008	3905	34259
2009	3946	32791

(b) Component means for the years 1996 to 2009 (balanced dataset).

Table 4: Estimated means (based on the log GDP data, but displayed on the original scale) of the three or respectively two distinct groups.

	low	middle	high
1970	0.50	0.33	0.17
1971	0.48	0.35	0.17
1972	0.43	0.40	0.17
1973	0.40	0.44	0.17
1974	0.39	0.44	0.17
1975	0.35	0.49	0.15
1976	0.30	0.56	0.14
1977	0.29	0.56	0.15
1978	0.31	0.54	0.16
1979	0.23	0.62	0.15
1980	0.33	0.51	0.16
1981	0.35	0.47	0.17
1982	0.37	0.46	0.17
1983	0.33	0.51	0.17
1984	0.33	0.49	0.18
1985	0.33	0.49	0.18
1986	0.32	0.50	0.18
1987	0.34	0.48	0.18
1988	0.30	0.51	0.18
1989	0.31	0.52	0.17
1990	0.31	0.50	0.19
1991	0.25	0.57	0.18
1992	0.42	0.36	0.22
1993	0.49	0.27	0.24
1994	0.50	0.25	0.24
1995	0.53	0.23	0.24

(a) Component weights for the years 1970 to 1995 (balanced dataset).

	low	high
1996	0.81	0.19
1997	0.81	0.19
1998	0.81	0.19
1999	0.81	0.19
2000	0.82	0.18
2001	0.81	0.19
2002	0.81	0.19
2003	0.82	0.18
2004	0.83	0.17
2005	0.82	0.18
2006	0.82	0.18
2007	0.83	0.17
2008	0.83	0.17
2009	0.83	0.17

(b) Component weights for the years 1996 to 2009 (balanced dataset).

Table 5: Estimated weights (based on the log GDP data) of the three or respectively two distinct groups.