

Desdoigts, Alain

Working Paper

Growth regression and counterfactual income dynamics

SFB 373 Discussion Paper, No. 2000,60

Provided in Cooperation with:

Collaborative Research Center 373: Quantification and Simulation of Economic Processes,
Humboldt University Berlin

Suggested Citation: Desdoigts, Alain (2000) : Growth regression and counterfactual income dynamics, SFB 373 Discussion Paper, No. 2000,60, Humboldt University of Berlin, Interdisciplinary Research Project 373: Quantification and Simulation of Economic Processes, Berlin, <https://nbn-resolving.de/urn:nbn:de:kobv:11-10047861>

This Version is available at:

<https://hdl.handle.net/10419/62184>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Growth Regression and Counterfactual Income Dynamics

Alain Desdoigts

On the one hand, it has been well established over the last decade that the worldwide income distribution has been polarizing or stratifying into distinct classes of income since World War II (see, among others, Bianchi (1997), Desdoigts (1994), and Quah (1996)). The analysis of the world income distribution is motivated by the following key question: can we find in the data evidences of poverty traps? Such a question can be found for instance in Baumol (1986)'s idea of convergence clubs.

On the other hand, an extensive empirical literature investigates the sources of growth and convergence through the estimation of worldwide cross country growth regressions using explicitly formulated growth models (see, among others, Mankiw, Romer, and Weil (1992), and Durlauf and Quah (1998) for an extensive and insightful survey on “why do growth rates differ?”).

As Quah (1996) stresses, growth regressions average across the cross-section but they can only give a picture of the behavior of the conditional mean, not of the whole distribution. This article illustrates how traditional cross-country growth regressions can be used to analyze the immediate sources of the worldwide income distribution dynamics over the period 1960–1985. As XploRe provides an extensive set of parametric and nonparametric methods, it is a natural statistical computing environment to investigate international growth differences and changes in the world income distribution.

The analysis starts with an estimation of a classical cross-country growth regression such as found by Temple (1998) who investigates the correlation between equipment investment and economic growth, and its compatibility with the Solow (1956) growth model from which an explicit convergence equation is derived. The estimated model and the underlying data are then used to quantify the immediate sources of each country's differential growth performance. In a second step, a nonparametric counterfactual exercise is proposed

that allows us to analyze the effects of the various explanatory variables on changes in the world income distribution. The effects of the different variables are estimated by applying kernel density methods. The procedure provides a visually clear representation of where in the density of incomes the specified factors exert the greatest impact. (See also Di Nardo, Fortin, and Lemieux (1996) for an analysis of the effects of institutional and labor market factors on the US distribution of wages, 1973–1992.)

1 A Linear Convergence Equation

```
{b, bse, bstan, bpval} = linreg(x, y)
  estimates coefficients for a linear regression problem from data x
  and y and calculates the ANOVA table
```

Following Mankiw, Romer, and Weil (1992), Temple (1998) estimates a linear **conditional convergence regression**, or **growth regression** using data for 78 countries and covering the period 1960–1985, that is of the form:

$$Y_i = \beta_0 + \beta_1 X_{i1} + \beta_2 X_{i2} + \dots + \beta_9 X_{i9} + \epsilon_i$$

where ϵ_i is a normally distributed error term reflecting a country-specific shock. The dependent variable is the log difference of output per working-age person over the period. The first four independent variables (X_1, \dots, X_4) are respectively the logarithm of average shares of real equipment and real nonequipment investment in real output, the logarithm of the average percentage of the working-age population that is in secondary school for the period 1960–85, and the logarithm of the annual average growth rate of the working-age population plus an exogenous rate of technological progress and a depreciation rate, both of them being constant across countries.

These variables reflect differences in factor accumulation across countries and are expected to control for growth differences in equilibrium. The fifth variable (X_5) is the logarithm of output per working-age person at the beginning of the period, and is expected to capture the Neoclassical convergence effect due to diminishing returns to reproducible factors, that tends to favor poorer countries. The last four exogenous variables (X_6, \dots, X_9) are dummies for respectively sub-Saharan Africa, Latin America and the Caribbean, East Asia, and the industrialized countries of the OECD plus Israel. These variables allow

us to control for differences in efficiency, variation of which has been found to be essentially intercontinental.

To estimate such a multiple linear regression, we first read the data Temple (1998) analyzed and that are stored in `temple` and define both the independent and the dependent variables.

```
z=read("temple.dat")
x=z[,2:10]
y=z[,1]
```

Second, we load the `stats` quantlib and use the following XploRe code that computes the linear regression of `y` on `x`, and stores the values of the estimated parameters as well as their respective standard error, *t*-statistic, and *p*-value.

```
library("stats")
{b,bse,bstan,bpval}=linreg (x,y)
```

This quantlet also provides as an output the following ANOVA (ANalysis Of VAriance) table that allows us to infer:

- that the model as it is specified allows us to explain about 80% of the variance exhibited in the annual average growth rate of income per working-age person over the period
- that the coefficient on the initial level of income per working-age person (X_5) is significantly negative; that is, there is a tendency for poor countries to grow faster on average than rich countries
- that the social returns to equipment investment in developing countries are very high (see also Temple 1998)
- that the variable reflecting the accumulation of human capital, in contrast to previous results, is not significantly different from 0, etc.

A N O V A	SS	df	MSS	F-test	P-value
-----	-----	-----	-----	-----	-----
Regression	10.957	9	1.217	33.397	0.0000
Residuals	2.479	68	0.036		
Total Variation	13.436	77	0.174		
Multiple R	= 0.90305				
R ²	= 0.81550				
Adjusted R ²	= 0.79108				
Standard Error	= 0.19093				

PARAMETERS	Beta	SE	StandB	t-test	P-value
-----	-----	-----	-----	-----	-----
b[0,]=	4.2059	0.7425	0.0000	5.664	0.0000
b[1,]=	0.2522	0.0354	0.5934	7.122	0.0000
b[2,]=	0.3448	0.0635	0.3966	5.426	0.0000
b[3,]=	0.0674	0.0533	0.1364	1.263	0.2108
b[4,]=	-0.4411	0.2480	-0.1476	-1.778	0.0798
b[5,]=	-0.3981	0.0543	-0.8488	-7.330	0.0000
b[6,]=	-0.2038	0.0828	-0.2178	-2.461	0.0164
b[7,]=	0.0642	0.0810	0.0676	0.793	0.4303
b[8,]=	0.3910	0.1175	0.2078	3.328	0.0014
b[9,]=	0.1611	0.1177	0.1747	1.368	0.1757

2 Counterfactual Income Dynamics

2.1 Sources of the Growth Differential With Respect to a Hypothetical Average Economy

Following De la Fuente (1995), we are now able to quantify the immediate determinants of growth and convergence during the period. The sources of the growth differential with respect to a hypothetical representative economy, basically the average country over the period, are computed using the above parameter estimates.

The following code decomposes each country's growth rate differential with respect to the sample average into five factors: the contribution of physical capital

accumulation, the impact of the working-age population growth, the contribution of human capital accumulation, the Neoclassical convergence effect, and the impact of a fixed effect reflecting differences in efficiency.

```
pkap=(x[,1]-mean(x[,1]))*b[2,]+((x[,2]-mean(x[,2]))*b[3,])
wagrowth=(x[,3]-mean(x[,3]))*b[4,]
hkap=(x[,4]-mean(x[,4]))*b[5,]
convergence=(x[,5]-mean(x[,5]))*b[6,]
fixed=((x[,6]-mean(x[,6]))*b[7,])+((x[,7]-mean(x[,7]))*b[8,])
      +((x[,8]-mean(x[,8]))*b[9,])+((x[,9]-mean(x[,9]))*b[10,])
```

2.2 Univariate Kernel Density Estimation and Bandwidth Selection

```
{hcrit, crit} = denbwsel(x{, h, K, d})
  starts an interactive tool for kernel density bandwidth selection
  using the WARPing method

fh = denest(x{, h, K, d})
  computes the kernel density estimate on a grid using the WARP-
  ing method

{fh, fh1, fhu} = denci(x{, h, alpha, K, d})
  computes the kernel density estimate and pointwise confidence
  intervals on a grid using the WARPing method
```

Suppose we are given a sample of independent, identically distributed realizations of a random variable $\{X_i\}_{i=1}^n$. Now, if a smooth kernel function $K\left(\frac{\bullet - X_i}{h}\right)$ is centered around each observation X_i and if we average over these functions in the observations, we obtain the kernel density estimate defined as follows

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

where the kernel function is a symmetric probability density function.

Practical application of kernel density estimation is crucially dependent on the choice of the smoothing parameter h . A measure of accuracy in order to assess

how closely $\hat{f}_h(x)$ estimates $f(x)$ is the Integrated Squared Error, $ISE(h) = \int (\hat{f}_h(x) - f(x))^2 dx$. Stone (1984) shows that a data-driven bandwidth \hat{h} that asymptotically minimizes $ISE(h)$ is given by

$$\hat{h} = \arg \min CV(h) \quad (2)$$

with $CV(h) = \int (\hat{f}_h(x))^2 - 2n^{-1} \sum_{i=1}^n \hat{f}_{h,i}(x_i)$ the cross validation function, and where $\hat{f}_{h,i}(x_i) = ((n-1)h)^{-1} \sum_{i \neq j} K\left(\frac{x-X_j}{h}\right)$.

Park and Turlach (1992) provide an overview over the existing bandwidth selection methods. We choose here to perform the Least Squares Cross Validation criterion instead of, for instance, the Biased CV or the Smoothed CV criteria that need either a very large sample size or pay with a large variance. Still, note that it remains difficult to recommend once and for all a particular bandwidth selector. One should therefore compare the resulting density estimates determined by different selection methods.

Our goal, here, is first to select an optimal bandwidth and second to estimate kernel densities of the world income distribution. We first load the necessary libraries. The `smoother` quantlib automatically loads the `xplore` and the `kernel` quantlibs. The `plot` quantlib is used for graphing the resulting cross validation and density functions.

```
library("smoother")
library("plot")
```

Second, we call the quantlet `denbwsel` that needs the univariate data vector as input and that will open a selection box which offers you the choice between different bandwidth selectors, as well as the possibility to change parameters such as the kernel, the search grid, etc. Among them, the LSCV criteria.

```
I60=x[,5] ./max(x[,5])
I85=(x[,5] .+y) ./max(x[,5] .+y)
{hcrit1,crit1}=denbwsel(I60)
{hcrit2,crit2}=denbwsel(I85)
```

Obviously, the CV function is not minimized within the automatically selected range of bandwidth. The bandwidth that minimizes the CV criterion is below the selected lower bound. We must increase the search grid for h . If one

manually selects a lower bound for h equal to 0.02, the following graphics are displayed that show the LSCV function in the upper left, the selected optimal bandwidth in the upper right, the resulting kernel density estimate in the lower left, and some information about the search grid and the kernel in the lower right. The graphical display is shown in Figure 1.

The optimal bandwidth corresponding to the world per working-age person income distribution in 1960 is therefore 0.039. It is stored in `hcrit1`. Note that the Sheather and Jones (1991)' selector chosen by Di Nardo, Fortin, and Lemieux (1996) finds a bandwidth equal to 0.038. We open a second selection box in order to compute the optimal bandwidth corresponding to the world per working-age person output in 1985. The lower bound of the search grid is now set to 0.015 and the corresponding optimal bandwidth obtained by least squares cross validation is now equal to 0.018. There is apparently more structure in the final distribution as compared to the initial distribution.

Confidence intervals can be derived under some restrictive assumptions (see Härdle 1991) and written as

$$\left[\hat{f}_h(x) - z_{1-\frac{\alpha}{2}} \sqrt{\frac{\hat{f}_h(x) \|K\|_2^2}{nh}}, \hat{f}_h(x) + z_{1-\frac{\alpha}{2}} \sqrt{\frac{\hat{f}_h(x) \|K\|_2^2}{nh}} \right] \quad (3)$$

where $z_{1-\frac{\alpha}{2}}$ is the $(1 - \frac{\alpha}{2})$ quantile of the standard normal distribution.

In XploRe, confidence intervals are computed using `denci`. The following quantlet code computes the confidence intervals for the optimal bandwidth previously selected by least squares cross validation, selecting a Gaussian kernel, a discretization binwidth d , and significance level $\alpha = 0.10$.

```
d=(max(I60)-min(I60))./200
{fh60,clo60,cup60}=denci(I60,hcrit1,0.10,"gau",d)
```

We propose now to decompose changes in the world income distribution on the basis of simple counterfactual densities. More specifically, and as proposed by De la Fuente (1995), what would the density of income have been in 1985 in a hypothetical world where the relative income of each country changed only due to factor accumulation, with all economies displaying average behavior in terms of all other variables? We simulate three such counterfactual densities. One is the density as defined above. Another is the density that one would have observed in 1985 if the relative income of each country changed only due

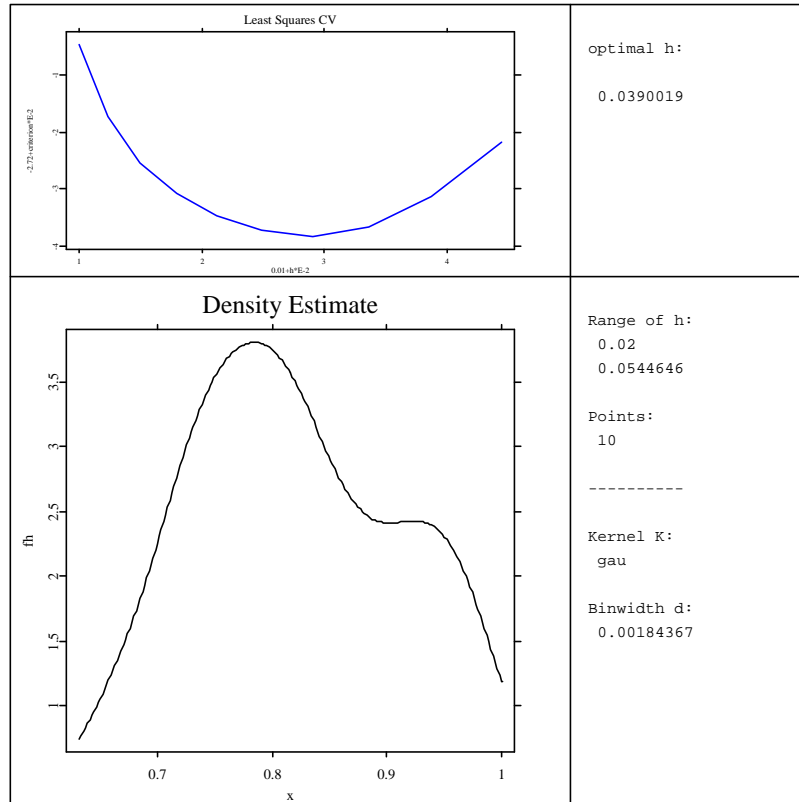



Figure 1: LSCV for the worldwide income per working-age person in 1960 normalized relative to the maximum.  `growdist.xlsx`

to the Neoclassical convergence effect. The last one is the density that the empirical model is able to predict.

In a first step, we compute the relative per working-age person income under the above assumptions and the observed density in 1985, and then estimate the corresponding counterfactual density. Comparing these densities with the density estimates corresponding to the real world in 1960 and 1985 gives a clear visual insight of the sources of the world income dynamics. Univariate density estimates are computed using `denest`. This quantlet only approximates the kernel density by the WARPing method. This method has the statistical efficiency of kernel methods while being computationally comparable to histogram methods as it performs smoothing operations on the bin counts rather than the raw data as in traditional kernel density estimation (see Härdle and Scott 1992). It is also possible to evaluate the density estimate at all observations by using `denxest` instead of `denest`.

```
w1=(x[,5].+pkap.+hkap.+wagr.+conv.+fix.+mean(y))
    ./max(x[,5].+pkap.+hkap.+wagr.+conv.+fix.+mean(y))
fhpred=denest(w1,hcrit2,"gau",d)
w2=(x[,5].+pkap.+hkap.+wagr.+mean(y))
    ./max(x[,5].+pkap.+hkap.+wagr.+mean(y))
fhfac=denest(w2,hcrit2,"gau",d)
w3=(x[,5].+conv.+mean(y))./max(x[,5].+conv.+mean(y))
fhconv=denest(w3,hcrit2,"gau",d)
fh85=denest(I85,hcrit2,"gau",d)
```

The above density estimates are displayed in Figure 2. To distinguish the densities, we choose to color them with the quantlet `setmask`. Technically, `setmask` handles mask vectors that contain numerical information to control the graphical display of the data points. This explains the name of the function. Density estimates are drawn as solid lines and confidence intervals as dashed lines.

```
fh60=setmask(fh60,"line","blue")
clo60=setmask(clo60,"line","blue","thin","dashed")
cup60=setmask(cup60,"line","blue","thin","dashed")
fh85=setmask(fh85,"line","red")
fhfac=setmask(fhfac,"line","yellow")
fhconv=setmask(fhconv,"line","green")
fhpred=setmask(fhpred,"line","magenta")
```




Figure 2: Univariate Density Estimates and Confidence Intervals. Upper left: Per working-age person income in 1960 (solid blue line) with pointwise confidence intervals (dashed blue lines) and in 1985 (red solid line). Upper right: Real (red line) and predicted (magenta line) per working-age person income densities in 1985. Lower left and right: Per working-age person income in 1960 (solid blue line) and counterfactual income densities in 1985 if countries would have differ only in factor accumulation (left) or in the Neoclassical convergence effect (right). [growdist.xpl](#)

To display Figure 2, we need to create a display which consists of four windows. This is achieved through the command `createdisplay`. The command `show` allows us to specify the data sets that will be plotted in each plot of the display. After `show` has been called, one controls the layout of the display by `setgopt`.

```

disp1=createdisplay(2,2)
show(disp1,1,1,fh60,clo60,cup60,fh85)
show(disp1,1,2,fh85,fhpred)
show(disp1,2,1,fh60,fhfac)
show(disp1,2,2,fh60,fhconv)
setgopt(disp1,1,1,"title","Density Confidence Intervals",
        "xlabel","Income in 1960 and 1985","ylabel",
        "density estimates")
setgopt(disp1,1,2,"title","Predicted vs 1985","xlabel",
        "Income (Predicted & 85)")
setgopt(disp1,2,1,"title","1960 plus factor accumulation
        effect","xlabel","Income (60 and factor accumulation)")}
setgopt(disp1,2,2,"title","1960 plus Neoclassical convergence
        effect",}
        "xlabel","Income (60 and convergence)")}

```

The upper left of Figure 2 displays both the density estimate of the per working-age person income in 1960 together with the corresponding confidence intervals (solid and dashed blue lines), and the per working-age person income density estimate in 1985 (red line). On the one hand, the distribution of income at the beginning of the period appears to be unimodal, most of the economies clustering in what one might call a middle-income class. On the other hand, the underlying density in 1985 seems to be consistent with a multimodal distribution suggesting that countries follow different development paths and that they tend to cluster into different income classes. The population of economies in 1985 seems to have at least three modes. The initial middle income class vanished: some countries caught up and joined a club of rich countries and others fell into a poverty trap. This is the “Twin Peaks” scenario illustrated among others by Quah (1996). At least, the structure of the worldwide income distribution in 1985 does not fit anymore within the computed confidence intervals corresponding to the income density estimate in 1960. There is a very systematic shift over times. This suggests a great amount of mobility within the system and over the period under study.

Where does this mobility exactly come from? What are the most important

factors in determining the worldwide income distribution dynamics? The upper right display of Figure 2 shows a counterfactual income density estimate that the empirical model estimated above has been able to predict together with the income density estimate in 1985. Although, the model appears to be able to predict the formation of the two modes for the highest income classes, and therefore to capture, at least partially, the convergence phenomenon, it is unable to fit the poverty trap that arose during the period. The lower left display suggests that differences in factor accumulation together with the differences in efficiency as proxied by the continental dummies, may be partially responsible for this wealth trap. But this cannot explain the all story. Something else is going on, and I leave here this issue for future exploration. Finally, the lower right display illustrates a collapsing over time of the world income distribution to a degenerate point limit. If all economies were displaying average behavior in terms of factor's accumulation and efficiency, poor countries would catch up with rich ones.

2.3 Multivariate Kernel Density Estimation

```
fh = denestp(x{, h, K, d})
      computes a multivariate density estimate on a grid using the
      WARPing method

gs = grcontour2(x, c{, col})
      generates a contour plot from a 3-dimensional data set x
```

All above formulas can be easily generalized to multivariate observations $\{X_i\}_{i=1}^n$, and $x_i = (x_{i1}, \dots, x_{id})^T$. The kernel function K has to be replaced by a multivariate kernel K^d . One takes a product kernel

$$K^d \left(\frac{x - X_i}{h} \right) = K^d \left(\frac{x_1 - X_{i1}}{h_1}, \dots, \frac{x_d - X_{id}}{h_d} \right) = \prod_{j=1}^d K \left(\frac{x_j - X_{ij}}{h_j} \right), \quad (4)$$

where $h = (h_1, \dots, h_d)$.

The following quantlet computes two-dimensional density estimates for different data sets via the function `denestp`, where the kernel is Gaussian and the bandwidth chosen arbitrarily to 0.05. The surface of each bivariate density


estimate is then illustrated via contour plots with contour lines $f(x, y) = c$. The function `grcontour2` allows us to generate contours corresponding to a bivariate density estimate.

```

library("graphic")
bi1=I60~I85
r=rows(bi1)
d=(max(bi1)-min(bi1))./20
fh1=denestp(bi1,0.05,"gau",d)
c1=(1:5).*max(fh1[,3])./10
gs1=grcontour2(fh1,c1)
bi2=I60~w1
d=(max(bi2)-min(bi2))./20
fhpred=denestp(bi2,0.05,"gau",d)
c2=(1:5).*max(fhpred[,3])./10
gspred=grcontour2(fhpred,c2)
bi3=I60~w2
d=(max(bi3)-min(bi2))./20
fhfac=denestp(bi3,0.05,"gau",d)
c3=(1:5).*max(fhfac[,3])./10
gsfac=grcontour2(fhfac,c3)
bi4=I60~w3
d=(max(bi4)-min(bi4))./20
fhconv=denestp(bi4,0.05,"gau",d)
c4=(1:5).*max(fhconv[,3])./10
gsconv=grcontour2(fhconv,c4)
disp2=createdisplay(2,2)
z=setmask(I60~I60,"line","red")
show(disp2,1,1,gs1,z)
show(disp2,1,2,gspred,z)
show(disp2,2,1,gsfac,z)
show(disp2,2,2,gsconv,z)
setgopt(disp2,1,1,"title","Bivariate Density Estimate",
        "xlabel","Income (1960)","ylabel","Income (1985)")
setgopt(disp2,1,2,"title","2D Density Estimate",
        "xlabel","Income (1960)",
        "ylabel","Predicted Income in 1985")
setgopt(disp2,2,1,"title","2D Density Estimate",
        "xlabel","Income (1960)","ylabel","Factor Accumulation")
setgopt(disp2,2,2,"title","2D Density Estimate",

```

```
"xlabel","Income (1960)","ylabel",  
"Neoclassical Convergence")
```

```
 growdist.xpl
```

The above quantlet leads to Figure 3. The display in the upper left box is the bivariate density estimate of per working-age person incomes in 1960 and 1985. If most observations concentrate along the 45°-line, then countries in the distribution remain where they started. In reality, poor (rich) countries do concentrate under (above) the 45°-line. Note also that whatever the class of income from which a country starts displays both catching up and lagging behind especially when a country started in the middle income class. This corroborates the emergent “twin peaks” in the cross-country distribution documented, for instance, by Quah (1996). This density estimate also corroborates the economic historian’s notion of convergence clubs; that is of countries catching up with one another but only within particular subgroups. If one isolates the Neoclassical convergence effect, then we obtain the density estimate displayed in the lower right box. Note how much the graph rotates counter-clockwise. This illustrates a potential for poor countries to overtake through the Neoclassical convergence effect. In fact, the twin peaks scenario arises mainly because of differences in the accumulation of reproducible factors (see the lower left display). However, the model as it is specified does not provide a perfect fit of the distribution dynamics at work under the period under study. In particular, it does not allow us to recover and to explain the formation of the poverty trap in the real distribution.

Still, counterfactual income dynamics as computed and analyzed above allow us to provide explanations to the regularities characterizing the evolution of the world income distribution. Although this is a new step in understanding cross country patterns of growth, much remains to be done. At least, this article provides an exercise which allows us to study the role of specific explanatory factors in explaining observed patterns of cross-country income distribution dynamics.

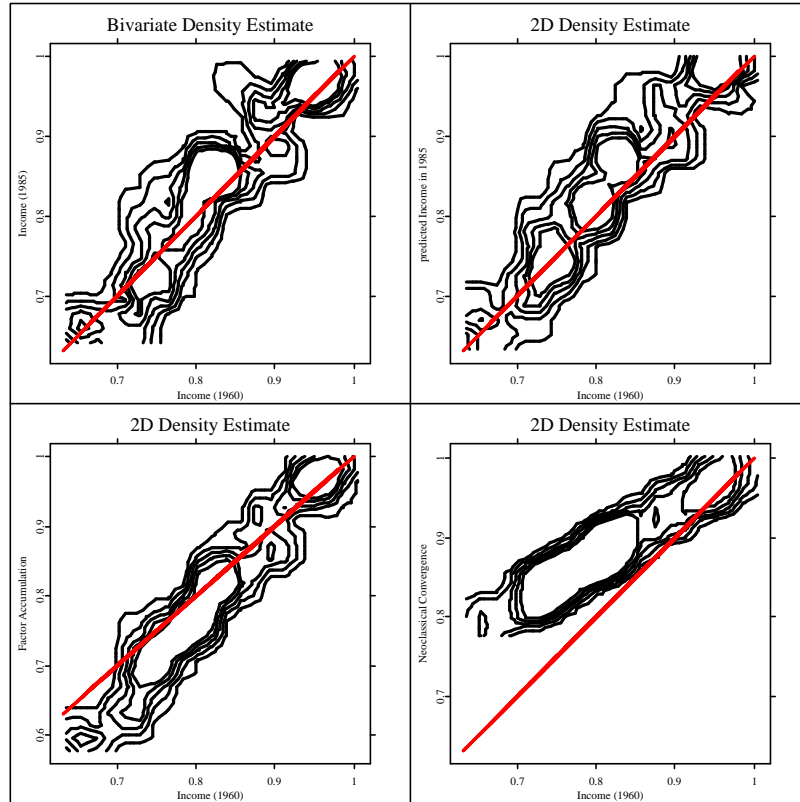


Figure 3: Contours of Bivariate Density Estimates. The x-axis is the per working-age person income in 1960. The y-axis is respectively: the per working-age person income in 1985 (upper left), the predicted income in 1985 (upper right), and the relative income of each country changed only due to factor accumulation (lower left) and to the Neoclassical convergence effect (lower right) with all economies displaying average behavior. [growdist.xpl](#)

References

- Baumol, W. (1986). Productivity growth, convergence and welfare: what the long run data show, *American Economic Review* **76**(5): 1072–1085.
- Bianchi, M. (1997). Testing for convergence: evidence from nonparametric multimodality tests, *Journal of Applied Econometrics* **12**(4): 393–409.
- De la Fuente, A. (1995). Catch-up, growth and convergence in the OECD, *CEPR Discussion Paper No. 1274* (November).
- Desdoigts, A. (1994). Changes in the world income distribution: A nonparametric approach to challenge the Neoclassical convergence argument, PhD. Thesis, European University Institute-Florence, June.
- Di Nardo, J., Fortin, N. M., and Lemieux, T. (1996). Labor market institutions and the distribution of wages, 1973–1992: a semiparametric approach, *Econometrica* **64**(5): 1001–1044.
- Durlauf, S. and Quah, D. T. (1998). The new empirics of growth, *Discussion Paper University of Wisconsin, Maddison and LSE* (January).
- Härdle, W. (1991). *Smoothing Techniques, With Implementations in S*, Springer, New York.
- Härdle, W. and Scott, D. (1992). Smoothing in by weight averaging using rounded points, *Computational Statistics* **7**: 97–128.
- Mankiw, N. G., Romer, D., and Weil, D. N. (1992). A Contribution to the empirics of economic growth, *Quarterly Journal of Economics* **107**(2): 407–437.
- Park, B. U. and Turlach, B. A. (1992). Practical performance of several data driven bandwidth selectors, *Computational Statistics* **7**: 251–270.
- Quah, D. T. (1996). Twin peaks: growth and convergence in models of distribution, *The Economic Journal* **106**: 1045–1055.
- Sheather, S. J. and Jones, M. C. (1991). A reliable data-based bandwidth selection method for kernel density estimation, *Journal of the Royal Statistical Society, Series B* **53**(3): 683–690.
- Stone, C. J. (1984). An asymptotically optimal window selection rule for kernel density estimates, *Annals of Statistics* **12**: 1285–1297.

- Solow, R. (1956). A Contribution to the theory of economic growth, *Quarterly Journal of Economics* **70**: 65–94.
- Temple, J. (1998). Equipment investment and the Solow model, *Oxford Economic Papers* **50**: 39–62.