

MPRA

Munich Personal RePEc Archive

Latin American clubs: uncovering patterns of convergence

Blyde, Juan

April 2006

Online at <http://mpa.ub.uni-muenchen.de/21320/>
MPRA Paper No. 21320, posted 16. March 2010 / 20:32

Latin-American clubs: uncovering patterns of convergence

Juan S. Blyde
Inter-American Development Bank

April, 2006

Abstract

This paper explores the dynamics of convergence in Latin American countries and asks whether there are tendencies for converging to different clubs. The analysis shows clear differences between two groups: a large group of low-to-middle income countries and a small group of rich countries. The club of low-to-middle income countries showed a tendency of spreading out until the mid 1990s and slight convergence afterwards. At the same time, the distance between the rich countries and the low-to-middle income countries faded away over time, particularly during the 1980s. However, during most of the 1990s, when convergence was occurring in the group of low-to-middle income countries, the rich countries started to pull away clearly distancing themselves again as a different club. The study of club behavior is important because the presence of clubs might suggest that there are common factors among groups of countries leading them to develop (and converge) in similar fashion. Identifying such common factors (if they exist) might improve our understanding of why some countries in the region grow faster than others. This is not possible to analyze with traditional growth regressions that employ a single catching up parameter or with a dispersion statistics like the sigma-convergence. Since these methods cannot detect club behavior much less they can analyze the reasons behind their formation.

JEL Classification: O40, O47

Key Words: Convergence clubs, polarization, stratification

Mailing address: Inter-American Development Bank, 1300 New York Avenue,
N.W. Washington, DC 20577. Stop E-1009. Tel: (202) 623-3517. Fax: (202) 623-2347.
E-mail: juanbl@iadb.org

1. Introduction

This paper uses a distribution dynamic approach to analyze the process of income convergence across Latin American countries. The approach has proved to be useful to study issues like the formation of clubs, polarization or stratification.

This paper departs from the traditional literature of economic growth that relies on econometric estimation to explain long-run growth. This literature normally encompasses a regression of the average growth rate of income against a set of conditioning variables and the initial income level. The idea underlying this procedure is that the conditioning variables explain the permanent growth component or trend, while the initial condition controls for transitory dynamics.

A significant drawback of this technique is that the transitory dynamics of the growth process are grossly over-simplified on a single parameter that seeks to measure the speed by which poorer countries catch up with richer ones. Unfortunately, a great deal of information is hidden in this parameter. For example, rich economies may be converging towards each other leaving the poor economies to form a different convergence club; or mid-income countries may be getting richer than those previously rich while the rich countries may be slipping towards the center of the distribution. Condensing all this information on a single parameter produces results that are often hard to interpret and, almost always, useless for policy advice.

Convergence regression and sigma-convergence analyses cannot reveal evolving distribution dynamics: they cannot show the mobility prospects in intra-distribution behavior, they cannot uncover the potential formation of clubs, stratification or polarization and they cannot foresee the limiting shape of the cross-section national income distribution. The distributional dynamic approach, however, is a technique especially developed to deal with these issues.

This paper explores the dynamics of convergence in Latin American countries and asks whether there are tendencies for converging to different clubs. One reason why this is relevant is that the potential formation of a club might suggest that there are common factors among a group of countries leading them to converge to a similar steady state. Identifying such common factors (if they exist) might improve our understanding of why some countries in the region grow faster than others. This is not possible to analyze with traditional growth regressions that employ a single catching up parameter or with a dispersion statistics like the sigma-convergence. Since these methods cannot detect club behavior much less they can analyze the reasons behind their formation.

The paper is divided in three sections besides this introduction. The second section looks at the evolving distribution of Latin American incomes and examines its dynamic regularities. Section 3 develops further dynamics, and examines intra-distribution mobility properties and longer-horizon tendencies. Section 4 summarizes the main findings.

2. The Shape of Distribution Dynamics

One convenient way to introduce the study of income dynamics is to look at a hypothetical income distribution. Figure 1, taken from Quah (1996a), plots on its horizontal axis the progress of time and on its vertical axis, per capita GDP. The figure shows at time t_0 a hypothetical income distribution across countries. As drawn, the distribution at t_0 has a single peak: most countries are in the middle of the distribution while only few are in the upper and lower tails.

There is one such distribution at each point in time. Figure 1 gives another one for t_1 . This latter distribution shows clustering at the high and low income levels while the middle incomes have almost vanished. Technically, the distribution at t_1 is said to show a bimodal shape and the sequence of distributions between t_0 and t_1 is known to exhibit an “*emerging twin peaks*” property.

Figure 1 also draws arrows that show intra-distribution transitions. Some countries with high incomes at t_1 had already high incomes at t_0 . Similarly, some countries with low incomes at t_1 began also poor at t_0 . However, there are other countries -indicated by the cross-crossing arrows in the figure- that transited from low to high income levels and conversely from high to low income levels. Therefore, this picture shows *persistence* and *mobility* simultaneously. Also the twin-peak shape shown in the figure carries the implication of *polarization* between rich and poor: the division between the rich, clustering together and the poor clustering together.

More than two peaks may also appear at t_1 . When this is the case, *stratification* is a more appropriate description than polarization. The intuition, however, does not change: distinct multi-tier income classes form and then separate from one another. They are also referred as *clubs*.

Cross-country relative income distribution

Figure 2 shows point-in-time snapshots of income distributions across 21 Latin American countries.¹ Per capita GDPs in each year are divided by the average per capita GDP across all the countries to obtain relative incomes. For each year, 1960, 1970, 1980, 1990, 2000, 2004, then the cross-country relative income distribution is found using a Gaussian kernel to estimate the density. The densities of these distributions are given in the figure with relative income values on the horizontal axis.

In 1960, the cross-country income distribution showed a significant peak at about 0.8 and two small shoulders at 1.8 and 2.4. In 1970, while the shoulder associated with the highest income is still visible, the shoulder in the middle of the distribution has disappeared. By this year the data shows that a formation of two distinctive clubs have emerged: a large club with an income around the region's average and a smaller club with an income more than two times higher.

By the beginning of the 1980s, however, things started to change. The higher income club moved closer to the lower income club on a trend that continues until the end of this decade. In fact, by the beginning of the 1990s, the two clubs stand very close to each other and the differences in the levels of the peaks are smaller too. This implies that the high income club is not as relatively rich as before and is also not as exclusive as before (in terms of the number of its members). The 1990s, however, seems to bring a reverse to the trends. By the year 2000, there is again a more clear distinction between the higher and lower income clubs in the sense that not only the distance between them has increased again but also the differences between their peak levels.

Tukey boxplot

A tool that helps us analyze further how the shape of the income distribution evolved over time is the Tukey boxplot shown in Figure 3. Time appears on the horizontal axis while the vertical axis maps income values. A single boxplot appears for each of the years. Consider the boxplot for 1960: coming down from above, the box itself begins at the top quartile (i.e., the 75th percentile) of the 1960 income distribution, and ends at the bottom quartile (i.e., the 25th percentile). Therefore, the height of the box is also the inter-quartile range which is a measure of spread in the distribution. Therefore, by construction the middle 50% of the

¹ Real per capita GDP is measured in purchasing power parity (PPP) . The data are obtained from the Penn World Table 6.1 for the period 1960-2000. The series have been expanded until the year 2004 using the growth rate of real GDP per capita from the World Development Indicators of the World Bank.

distribution lies entirely within the box. If the inter-quartile is large –the box is tall– the middle 50% of the distribution is relatively spread out. By contrast, if the inter-quartile is small –the box is short– the middle 50% is tightly compressed together.

The horizontal bar that appears somewhere inside the box is the median of the income distribution. The upward and downwards distance from the median to the top and bottom of the box provide information of the shape of the income distribution. If these distances differ, then the distribution is skewed or asymmetric.

Thin vertical lines emanate from the box both upwards and downwards: these end at what are known as upper and lower adjacent values, respectively. If the inter-quartile range is r , then the upper adjacent value is the largest income value observed that is no greater than the top quartile plus $1.5 \times r$. Similarly, the lower adjacent value is the lowest income observed that is no smaller than the bottom quartile minus $1.5 \times r$.

Finally, there are outliers or observations outside the upper and lower adjacent values. They are individually indicated with asterisks. Note that upper and lower outside values might not exist –the adjacent values might already be the extreme points in a particular realization.

The Tukey boxplot shows that the distribution had a general tendency of spreading out almost during the entire period. During the 1960s and 1970s, for example, this spreading out occurs as a consequence of the 75th percentile growing faster than the 25th percentile. During the crisis of the 1980s, however, the 25th percentile did not grow anymore and in fact, together with the median, started to decrease. The 75th percentile, however, kept growing contributing to the dispersion. The 1990s, saw a recovery not only of the growth in the 25th percentile but also of the median but, again, at least until 1995, the 75th percentile was growing faster and the inter-quartile range increased even more. It was only after 1995 that the 75th quartile stopped rising and fell slightly. This actually prompted a reduction in the dispersion of the middle part of the distribution. The general tendency over the four decades, however, was that the middle part of the distribution spread out with the maximum occurring sometime around 1995.

What happened to the upper part of the distribution (the upper adjacent values and the outliers) is interesting. As the inter-quartile range increased so do the upper adjacent values. Therefore, some outliers were covered by the rise of the upper adjacent values and others fell in their positions. Between 1985 and 1990, however, the upper adjacent value decreased. This is consistent with what we saw in figure 2 that by 1990, the separation between the rich club and the rest of the distribution was the smallest. The 1990s, however, put an end to this trend. By

1995, the upper extreme point (although still within the upper adjacent value) increases dramatically, and by 2004 there were already two outliers outside the upper adjacent values again. Therefore, after the crises of the 1980s, there was a push again for a more clear separation between the high income club and the rest of the distribution as we also saw in figure 2.

We can summarize these findings as follows: while the middle part of the distribution became more spread out over four decades, the upper part of the distribution was pulled towards the center, particularly during the 1980s, and then pushed back since the beginning of the 1990s.

It is useful to study the Tukey box with the numerical values shown in table 1. This table shows statistics that match certain features in the figure. The increase in the inter-quartile range until 1995 and the slight reduction afterwards is clearly seen in the table. It is interesting to note that while the inter-quartile range was increasing, the standard deviation showed a general decreasing trend particularly during the 1980s. This reveals that the distribution was actually getting less spread out in its tails even though it was getting more spread out in the middle. Again, this is consistent with the story of the higher income club was being pulled towards the rest of the distribution during the bad times of the 1980s. After the 1980s, however, the growth of the inter-quartile range first slowed down and then became negative while the standard deviation kept growing. This indicates that the spreading out of the distribution at this time occurred in the tails (particularly in the upper tail) and not in the middle, another indication that the club of rich countries was pushing itself away from the rest of the distribution.

In this section we have provided some preliminary evidence of difference convergence behavior in Latin America: first, there seems to be a small group of rich countries differentiating themselves from the group of low-to-middle income countries. The income of the low-to-middle income countries in general showed a tendency of spreading out at least until the mid 1990s. Then their incomes showed a slight pattern of convergence. At the same time, the distance between the rich countries and the low-to-middle income countries was fading away over time, particularly during the 1980s. However, during most of the 1990s, when convergence was occurring in the group of low-to-middle income countries, the rich countries started to pull away clearly distancing themselves again as a different club. Of course, the countries of the rich club at the beginning of the sample do not have to be the same as those at the end of the sample. To study that systematically, one needs to analyze intra-distribution properties, the topic of the next section.

3. Mobility and dynamics

Thus far we have discussed only point-in-time snapshots of the distribution. This section takes the next step: it investigates intra-distribution dynamics and long-run tendencies. In terms of figure 1, this section considers two features there: first, the transitional arrows in bold, and second, t_1 being outside the sample, possibly arbitrarily far outside.

Cross-profile dynamics

We start with a figure of cross-profile dynamics.² Figure 4 illustrates when economies overtake one another, fall behind, or pull ahead. The vertical axis is the log of relative incomes.

Each curve in the figure shows the situation at a given point in time: 1960, 1975, 1990 and 2004. The lowest curve is that for 1960; proceeding upwards, we find curves for progressively later years. The lowest curve, moreover, is constructed to be monotone rising: at the earliest time period, the cross section of economies is sorted in increasing order. This ordering is then maintained throughout the time periods considered. Obviously, the character of the upper plots depends on the initial periods when the ordering is taken. The change in character of the plots through time, however, does not.

If, over time, each cross profile always monotonically increases, then income ranking are invariant. However, if the cross-profile plots exhibit choppiness this is an indication that intra-distribution churning is occurring. Additionally, the slope of the cross-profile plot shows the income distance across the rich and the poor.

Table 2 contains numerical measures of the slope features in figure 4. In the first column, Slope measures how much, on average, (the log of relative) incomes rise as one proceeds rightwards across economies in a given cross-profile curve. If Slope has value x , then “being 10 economies richer” means having income e^{10x} time higher.

The table shows that the Slope decreased over time, meaning that the differences in income across countries were less pronounced (in general) particularly from 1960 to 1990 when the reduction of the slope was more dramatic. The same can be said by looking at the second column. In fact, the first two columns only provide support to our previous finding that the upper tail of the distribution approached the inter-quartile range particularly until the end of the 1980s. This is the main

² This entire section was done using Quah’s TSRF econometrics shell.

reason why the differences in income across countries were becoming less pronounced during this period.

The increase in choppiness through time in the cross-profile plots is an indication of the lack of stability in the relative positions of the countries over time. By 2004, local peaks -around economies ranked 2nd, 6th, 9th, 17th poorest in 1960- had appeared: these turned out to be: Dominican Republic, Panama, Brazil and Chile. By contrast, Honduras (3rd poorest in 1960), Bolivia (8th), Nicaragua (13th) and Venezuela (21st), had experienced economically significant relative declines by 2004.

The cross-profile dynamics in figure 4 and table 2 have given a first, preliminary look at the importance of figure 1's transition arrows: they, however, do not yet show any deep underlying regularities in the data. I turn therefore to two more representations of intra-distribution churning: first, transition probability matrices, and second stochastic kernels.

Transitional probability matrix

We now track the evolution of the entire income distribution over time using Markov chains. Markov chains are employed to approximate and estimate the laws of motion of the evolving distribution. The intra-distribution dynamics information is normally encoded in a transition probability matrix.

The transition probability matrix in table 3 shows 40-year transitions between the distributions of relative GDP per capita between 1960 and 2004.³ The matrix is partitioned in 5 brackets of equal length in relative income (1: 0-0.5, 2: 0.5-1.0, 3: 1.0-1.5, 4: 1.5-2.0, 5: 2.0 and higher) . The first column of the table (labeled under *data*) shows the percentage of the data in each bracket at the beginning of the 1960s. Clearly, most of the countries belonged to brackets 2 and 3 at the beginning of this decade, which is also evident from figure 2 (first panel).

The 45-degree diagonal shows the proportion of countries that remain in the same bracket of relative income after 40 years. The last row, for example, shows that for the countries that exhibited a GDP per capita below 0.5 times LAC's average at the start of the 1960s, 89% remained in the same bracket 40 years later, while 11% experienced an increase in their relative incomes to the 1-1.5 bracket. None of the countries moved higher up in the distribution.⁴

³ The number are actually the average of 5 transitions: 1960-2000, 1961-2001, 1962-2002, 1963-2003, 1964-2004.

⁴ Although we use plural here, some of these percentages might be associated with only one country.

The second row (from the bottom) shows that within the bracket 0.5-1.0, persistence was lower: 57% of the countries stayed in the same bracket while 11% moved down to the lowest bracket. 32% of the countries, however, moved up one bracket in the distribution. Within this income range there was a larger tendency to move up in the distribution than to move down. Persistence was also low in the 1.0-1.5 bracket. Only 17% of the countries that were in this bracket at the beginning of the 1960s were found in the same bracket after 40 years. So mobility was also the norm here: 42% of the countries fell to the bracket immediately below and even 4% fell two brackets below. Only 17% of the countries moved up in the distribution showing that within this range there was a larger tendency to move down in the distribution than to move up. In the range of countries with incomes between 1.5-2.0 times LAC's average, persistence was again very high with 83% of the countries keeping their relative positions. After the four decades, 17% of the countries in this group moved up. Finally, for the countries that started with a relative income higher than 2.0, only 10% remained in this range. 40% of the countries moved down to the 1.5-2.0 range and 50% moved down to the 1.0-1.5 range.

Stochastic kernel

One shortcoming of the transition probability matrix is that the selection of income brackets is arbitrary –different sets of discretisations may lead to different results. One way to overcome this drawback is by using the stochastic kernel.

The stochastic kernel replaces the discrete income states of the transition probability matrix by a continuum of states. This means that we no longer have a grid of fixed relative income states, like [0.0-0.5), [0.5-1.0), etc, but allow the states to be all possible intervals of income. By this we remove the arbitrariness in the discretisation of the brackets. We now have an infinite number of rows and columns replacing the transition probability matrix. The stochastic kernel can be viewed as a continuous version of the transition probability matrix. Figure 5 shows the stochastic kernel for relative per capita income of 40-year transitions between 1960 and 2004.

Interpretation of the stochastic kernel is as follows. Stand at any point on the Period t axis, and then look in a straight line parallel to the Period $t+40$ axis. What is traced out on the surface of the kernel is a probability density: it is always non-negative and integrates to 1. The more likely transition possibilities manifest as higher values on this line. The 45-degree diagonal indicates persistence properties: regions in different parts of the income distribution remain roughly where they begin. The more pronounced the surface along this diagonal, the more persistence there is in the distribution. The lower is this diagonal, the greater is the intra-distribution mixing. A single ridge parallel to the Period t axis indicates

convergence: poor regions growing faster and rich regions slowing down so eventually all income levels are equalized. Another extreme occurs when piling up takes place along the negative-sloped diagonal which shows regions dynamically overtaking one another. Peaks indicate club behavior: regions around a particular income class become attached, over time, to precisely such income class (this is a peak that is on the 45-degree diagonal) or to a different income class (this is a peak that is off the 45-degree diagonal).

Observation of the stochastic kernel in figure 5 shows the presence of two hills. The hill closer to the lower income is less pronounced indicating that more intra-distribution mixing occurred in this area while relatively more persistence happened at the higher incomes.⁵

The hill closer to the lower incomes indeed resembles a ridge. This ridge lays somewhat close to the 45-degree diagonal. For certain income ranges, however, the ridge is slightly off the diagonal. For example, for very low income levels the ridge is situated slightly below the diagonal and for mid income levels the ridge lays slightly above the diagonal. This can be verified in the contour plot of the stochastic kernel which is shown in figure 6. For regions with incomes around 0.8 times the LAC's average (and below), the highest probability occurs somewhat at higher incomes, so these regions had a higher tendency to move slightly up in the distribution than to a lower part. Countries with income levels between 0.8 and 1.5 times, however, experienced a lot of intra-distributional mixing: several countries saw an improvement in their relative position but several others saw a deterioration. However, the probability to move down in the distribution was slightly higher as the ridge is somewhat above the diagonal. This is consistent with the results from the transition probability matrix. The findings in this range of incomes indicate the existence of some signs of cohesion around an income of 0.8 the LAC's average. Note that this is not inconsistent with the results from section 2 that indicate that the middle part of the distribution showed a general tendency of spreading out almost during the entire period. While some cohesion was being formed around the 0.8 range, the distribution in the middle became more disperse mostly because of the many countries between the 0.8 and 1.5 range improving their relative positions as well.

Some of the countries within the 0.8-1.5 range that improve their relative position over the course of 40 years actually became part of the rich club at the end of this period. The rich club is seen in the contour plot around the 1.7 relative income.

⁵ It is worth mentioning that the hill closer to the lower incomes has more grid lines drawn than the other hill and so the resolution is correspondingly finer. The reason for this is that this is the place where more observations are available. The hill closer to the higher incomes has relatively fewer observations. The precision of the estimation is lower in this area.

Note again that the contour plot is not tight around this income which means that some distributional mixing occurred here as well: some countries coming from below eventually joined this club and some countries originally members of this club felt down in their position.

It is instructive to identify which countries belong to which club and to track the evolution of their membership over time. This is done in table 4. The table divides the sample into two clubs corresponding roughly to what we found in the stochastic kernels. Club 1 encompasses countries with incomes between 0 and 1.5 times the LAC's average and Club 2 encompasses countries with incomes 1.5 and higher.

The first row shows the countries that kept their relative positions after four decades. These are the countries whose relative income's trend did not increase or decrease more than 10 percent from the beginning to the end of the sample period. The second column shows the countries that moved up in their relative position but stayed in the same club. The third column shows the countries whose position deteriorated but stayed in the same club. Finally, the fourth and fifth columns show the countries that changed memberships: countries that moved from club 1 to club 2 (fourth column) and countries that moved from club 2 to club 1 (fifth column). The complete evolution of each country is shown in figures 7.1 to 7.21.

Argentina is a country whose relative position deteriorated over time, however, the country started with such a high relative income at the beginning of the 1960s that the fall was not enough to change its club membership. That was not the case for Venezuela, which started with an even higher income than Argentina but by the early 1990s had already switched clubs. Chile and Mexico had the opposite experience. They both switched memberships from club 1 to club 2. Brazil showed a tremendous increase in its relative position across the Latin American sample, however, the country started with a much lower level than Chile or Mexico and so the improvement was not enough to make the switch to club 2. Something similar happened with the Dominican Republic. Other increases in relative positions -but not enough to switch clubs- occurred in Colombia, Ecuador, Haiti, Panama and Paraguay. Countries that showed persistence in their relative positions were Costa Rica, Guatemala, Guyana and Haiti.

Bolivia, Peru, Honduras, Jamaica, Nicaragua and El Salvador, all members of club 1, saw deteriorations in their positions. It is interesting, however, that out of 19 countries that started in club 1, only 6 of them saw their position deteriorated. This confirms our earlier findings that across the low-to-middle income countries, there was a relatively lower tendency to move down in the distribution than to stay in the same place or to move up.

Finally, Uruguay, a country that started in the rich club, not only managed to stay within the same club but also to improve its relative position.

So far we have analyzed dynamics in the past. However, the distribution dynamic approach can also be used to explore the shape of the distribution in the future. In order to do this, we go back to the transition probability matrix. We can assume that the system described in table 3 continues in its evolution. In other words, the transition that took place between 1960 and 2004 occurs again several times. If nothing structural were to change, then the system could eventually reach a long-run steady state. This is what is called the ergodic distribution (or the steady state distribution), and it is shown in the last row of the table. This is obtained by iterating the transition probability matrix repeatedly over the long horizon to get an estimation of the long run tendency of a country to land up in a given income bracket.⁶

The long run tendency characterized by the ergodic distribution has two distinctive characteristics: first, the distribution is more spread out relative to the one in 1960s. Second, the peak is more biased to the upper tail. Our earlier results help to explain this features: first, we already saw that across the large group of low to middle income countries there was a larger tendency to move up in the distribution than to move down. This implies that the peak of the distribution was slowly moving up to larger relative incomes. However, in this process, still some countries maintained their relative positions and some even lost them. Therefore, while the peak was slowly increasing towards higher relative incomes, its height was slowly falling because of the countries that were not able to keep up with the increase. This is the reason why the distribution is more spread out in the low to middle income range.

4. Concluding remarks

In this paper we have uncovered the presence of club behavior within Latin American countries. Countries have shown patterns of convergence leading them to join different clubs. The analysis shows clear differences between two groups: a large group of low-to-middle income countries and a small group of rich countries. The club of low-to-middle income countries showed a tendency of spreading out until the mid 1990s and slight convergence afterwards. At the same time, the distance between the rich countries and the low-to-middle income

⁶ It should be mentioned that the ergodic distribution shall not be read as a forecast of the future. Government policies might change, important events might occur, and many forces behind the widening of income gaps might exhibit non-linear effects. Therefore, the ergodic distribution shall be interpreted simply as the long run realization of similar tendencies in the absence of structural changes (Quah, 1996a).

countries faded away over time, particularly during the 1980s. However, during most of the 1990s, when convergence was occurring in the group of low-to-middle income countries, the rich countries started to pull away clearly distancing themselves again as a distinct club.

Uncovering club behavior, however, is the first and probably easiest step for understanding the development processes behind these countries. More difficult would be to extract policy implications from these patterns.

For example, are there common factors behind the fact that Chile and Mexico were in the poor-to-middle income club -with relatively similar income levels in the 1960s- and both were able to advance to the rich club?, or this was the result of completely different factors?. Likewise, are there common factors behind the fact that Brazil, Colombia and Panama were in the poor-to-middle income club -with relatively similar income levels in the 1960s- but, unlike Chile and Mexico, were not able to advance as significantly? Are there common factors behind the fact that Bolivia, Peru, Honduras, Jamaica, Nicaragua and El Salvador had relatively poor incomes in the 1960s and all became relatively poorer after 40 years?

This paper has departed from the traditional literature of economic growth that relies on econometric estimation to explain long-run growth. A significant drawback of this technique is that the transitory dynamics of the growth process are grossly over-simplified on a single parameter that seeks to measure the speed by which poorer countries catch up with richer ones. The methodology used in this paper has allowed us to uncover different patterns of convergence among groups of countries leading them to form different convergence clubs.

Club behavior also implies that there may be common factors among groups of countries. Whether similar convergence patterns are the result of common factors or unrelated events should be the topic of future research.

References

- [1] Barro, R. and X. Sala-i-Martin (1991) "Convergence across states and regions", *Brookings Papers on Economic Activity*, 1.
- [2] Barro, R. and X. Sala-i-Martin (1992) "Convergence", *Journal of Political Economy*, 100.
- [3] Ciccone, A. and R. Hall (1996) "Productivity and the density of economic activity" *American Economic Review*, 67.
- [4] Cowell, F. (1995), *Measuring Inequality*, Harvester Wheatsheaf, London
- [5] Duro, J.A. and J. Esteban (1998) "Factor decomposition of cross-country income inequality", *Economic Letters*, 60.
- [6] Easterly, W. (2005) "National Policies and Economic Growth" in P. Aghion and S. Durlauf, eds., *Handbook of Economic Growth*, vol. 1A, North-Holland, Amsterdam, 2005.
- [7] Robinson, S. (1976) "A note on the U-hypothesis relating income inequality and economic development", *American Economic Review*, 66.
- [8] Theil, H. and C. Moss (1999) "The measurement of inequality by components of total expenditure" *Empirical Economics*, 24.
- [9] Quah, D.T. (1996a) "Convergence empirics across economies with (some) capital mobility", *Journal of Economic Growth*, 1.
- [10] ----- (1996b) "Twin peaks: growth and convergence in models of distribution dynamics" *Economic Journal*, 106.