12. The Hump-Shaped Transition Path for the Growth Rate

The chapters of Part II of the book have showed that a handful of socioeconomic variables have underlying transition paths. Many other socioeconomic variables have such paths. This chapter demonstrates that the aggregate GDP data has a transition too. Development is a process that diverges from the traditional steady state and converges to the modern one. Thus, the growth rate moves from zero to about two. If countries catch up, there must be, at least, one peak in between. The chapter shows that this is indeed the case, and hence countries do catch up.

The eight sections of Chapter 12 start by presenting the data (s1). Then the hump-shape is estimated (s2), and its robustness is demonstrated (s3). Next, the theoretical consequences are discussed – the key point is that the form found rejects the standard one-sector model (s4), but it is easy to explain by the good old two-sector model with a traditional sector that is gradually replaced by a modern sector (s5-6).[1] Finally, the hump-shape is replicated with non-linear panel regressions; the form is significant, but the explanatory power of the shape is small (s7).

12.1 Data and four samples: All, Basic, Main and OPEC

This chapter uses GDP-data only. As before, the real per capita GDP is the *gdp*. It is the *cgdppc* data from the Maddison Project. The *gdp* is used to calculate the growth rate, $g = \Delta gdp/gdp_{-1}$, and income, $y = \ln gdp$. These data are used to estimate the relation between the annual growth rate of a country, g_{it}, and the income level, $y_{it(-)}$, at the beginning of period *t*. The countries included account for more than 95% of the world population. As no other series is used in the chapter, the wide sample goes back to 1950, where about 60 countries, notably in Africa, were still colonies. Thus, the time dimension is 1950-2016.

The sample of *All* data is reduced to the *Basic* sample by a deletion of outliers, defined as all observations in the first and in the 99th percentiles of the growth rates and the income levels; see Figure 1. This procedure discards observations of growth rates above 26% and below -21% and *gdp* levels below \$610 and above \$57,250, where *income* is 6.4 and 11, respectively. As before, the Basic sample is divided into the *Main* and the *OPEC* samples. Table 1 gives some descriptive statistics for these samples.

[1] In Chapter 10 *T* was used for the corruption index (from Transparency International), but in the present chapter *T* is used for the traditional sector.

Table 1. Descriptive statistics for alternative samples, 1950-2016

Sample name	ALL (10,329)			Basic (9,931)			Main (9,137)		
Deleted	None			Extreme (398)			OPEC (794)		
	gdp_{it-1}	y_{it-1}	g_{it} %	gdp_{it-1}	y_{it-1}	g_{it} %	gdp_{it-1}	y_{it-1}	g_{it} %
Mean	9,401	9.15	2.5	8,666	9.07	2.5	8,563	9.06	2.5
Standard dev.	13,292	9.49	8.2	9,971	9.21	6.4	9,935	9.20	5.8
Maximum	220,71	12.30	133.6	56,319	10.94	26.0	56,319	10.94	26.0
Minimum	134	4.90	−62.9	609	6.41	−21.0	609	6.41	−21.0
1st percentile	608	6.41	−21.0	701	6.55	−14.8	695	6.54	−14.5
50th percentile	4,592	8.43	2.5	4,580	8.43	2.5	4,381	8.39	2.5
99th percentile	57,244	10.96	26.0	44,308	10.70	19.6	44,281	10.70	18.9
Years	66			66			66		
Max *N*	169			168			153		
Min *N* per year	139			130			119		

The real GDP per capita in 2011 US$ is termed ***gdp***. It is the *cgdppc* series from Maddison Project (2018). ***Income***, *y*, is the logarithm to *gdp*. ***Growth***, *g*, is for *gdp* in year *t* compared to year *t*-1. Extreme observations are below the 1st percentiles and above the 99th.

Extreme negative growth rates often reflect (civil) wars, the breakdown of the Soviet Union, or large negative changes in oil prices. Extreme positive growth rates are due to the exploitation of newly discovered natural resources and mean reversion after large negative shocks. A *gdp* level of $610 is compatible with the lowest *gdp* levels that are recorded in the Maddison Project Database for pre-industrial times: the first percentile of *gdp* is at $650 (where y = 6.5) for all countries and all years before 1750. Hence, still lower income levels are extreme by historical comparison. Income levels beyond the 99th percentile are also extreme. Extreme high-income observations are dominated by OPEC countries with a small population.

12.2 Estimating the hump, the best estimate of the transition curve, $g = \Pi^g(y_{(-)})$

Figure 1 shows a scatter plot of the correlation between the growth rate and the income level. Each dot represents one of the 10,329 country-year observations for pairs of *g* and initial $y_{(-)}$ for 1950-2016. The dashed vertical and horizontal lines identify observations above and below the first and the 99th percentiles of the growth rate and the income level. Gray dots represent observations for OPEC countries, which account for a large fraction of the outliers. The inner rectangle gives the observations for the Basic sample, where outliers have been deleted.

The wild scatter of the data points and the packed rectangle mean that any presumed growth path can only explain a small fraction of the variation at best. The thick black line through the middle of the inner scatter is the ***best*** kernel regression estimate of the transition-path, Π^g, for the sample without outliers (Basic). With a range of the vertical axis (annual growth rate) from +125% to −75%, the reported kernel line looks flat.

Figure 1. The growth-income scatter for the full sample, 1950-2016

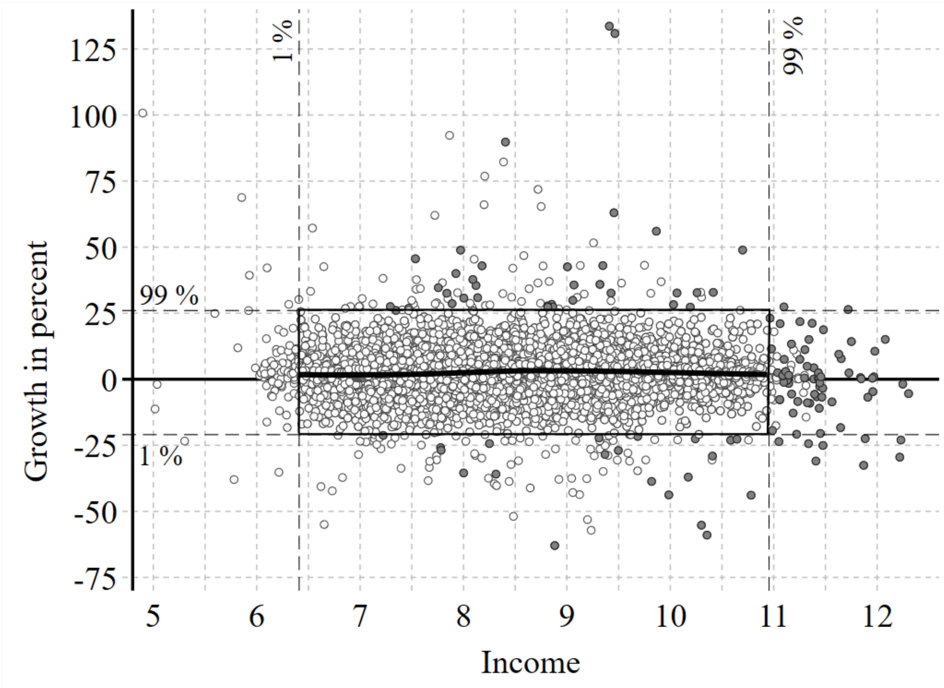

The inner square is the Basic sample used for the estimated kernel of Figure 2. See Table 1. The gray observations are from the OPEC countries excluded in the Main sample.

Figure 2. The enlarged growth-income path for the Basic and Main samples

The black curve is taken as the best estimate of the transition curve $g = \Pi^g(y_{(-)})$

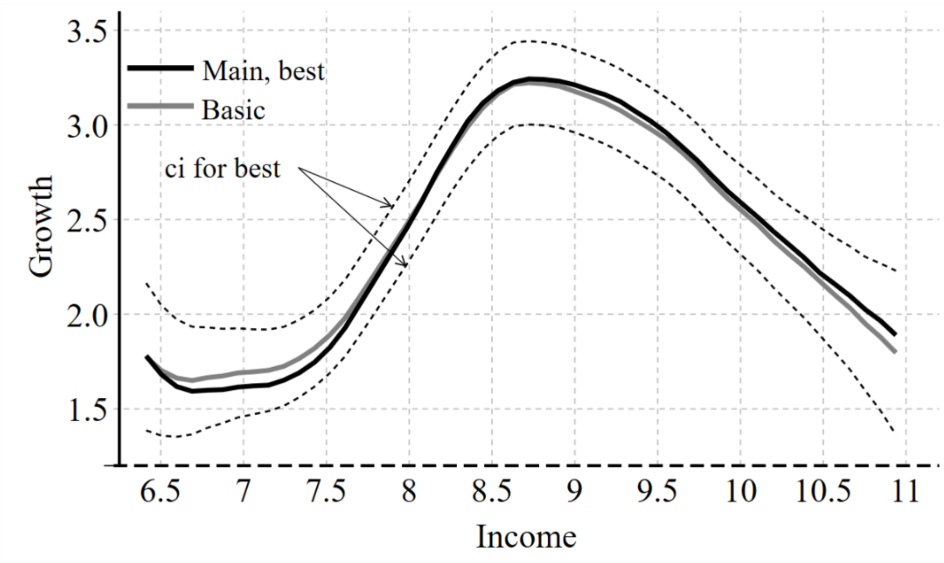

Both curves are estimated for the bandwidth chosen by Stata, which is 0.32 and 0.31 for the two samples. It gives virtually the same curves to use *bw* = 0.3 in both cases. The 95% confidence intervals are drawn for the Main sample. The curves are within the confidence intervals of each other. As usual, the 95% confidence intervals (ci) are rather narrow thanks to the large number of observations. All curves within the confidence intervals are hump-shaped, but the peak is only determined within the interval from y = 8.2 to 9.7, and for g = 3 to 3.4.

Figure 2 gives an enlarged picture of the same growth-income path, where the scatter points are suppressed. Zooming in reveals a hump-shaped growth path. It picks up at an income level of about 7.3, peaks with a growth rate above 3% at an income level of about 8.7, and thereafter falls toward a potential steady state growth rate slightly below 2%. Growth at the low end is 1½% and thus well above the historical growth rates before modern growth started. This means that the transition has started even in the poorest countries as argued in Chapter 1.8. However, the path at the low end (from 6.3 to 7.3) is flat. There seems to be no low-level equilibrium trap. It is not easy to start growing, but Malthus' mechanism is not evident; see also Figure 1.1b at the start of the book.

The graph also shows that it does not matter for the estimated growth path whether OPEC members remain in the sample or not as long as extreme observations are deleted. There is an almost perfect overlap between the paths estimated on the Basic sample (black line) and the Main sample (gray line). The estimated rule-of-thumb bandwidths are both close to 0.3 (used from now), and the confidence intervals are so similar that only one is shown. The best estimate of the transition curve $\Pi^g(y_{(-)})$ is estimated on the slightly less noisy Main sample (no outliers, no OPEC countries). The reported confidence interval is sufficiently narrow to rule out any line where the slope has the same sign over the full income range. Thus, the workhorse model does not work for the full range. Hence, the main empirical result is that the growth-income path is hump-shaped, with a sign change in the middle.

This means that from an income level of 7.5, low-income countries do catch up with high-income countries. However, the growth of the average less developed country peaks at 3.25%, which gives a catch up rate of only 1.5%, so the time necessary to close a gap of say 25 times is a bit more than 200 years.

12.3 Robustness of the best Π^g-curve from Figure 2

As in the previous chapters, a set of experiments are made to show that the transition curve is robust. Figure 2 uses the annual data – most growth studies use a longer time unit, such as 5- or 10-year averages. Figure 3a shows that it matters little for the form of the curve. The peak is a little lower for a longer time unit, and as N falls, the confidence interval widens.

Figure 3b demonstrates how alternative bandwidths affect the estimated kernel line. The kernel line becomes quite wobbly for a low bandwidth of 0.1 and approaches a straight line for a high bandwidth of 0.9. In between, the hump-shaped transition path remains. The kernel line for bw = 0.1 to 0.5 is almost completely within the confidence interval for the reference kernel.

Figure 3. Robustness of the best *Πg-curve* from Figure 2, shown as the bolded black curve

Figure 3a. Annual, 5-year and 10-year time unit

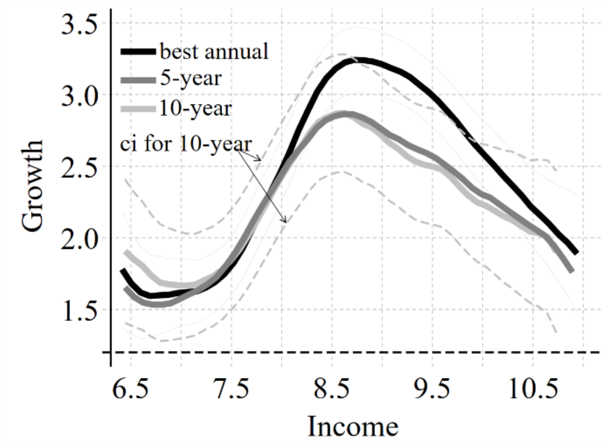

Confidence intervals are for the 10-year unit

Figure 3b. Experiments with bandwidth

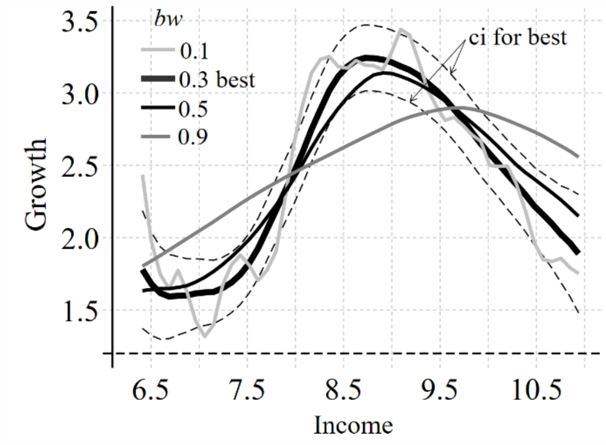

Confidence intervals are for the $\Pi^g(y, 0.3)$-curve

Figure 3c. Excluding time periods

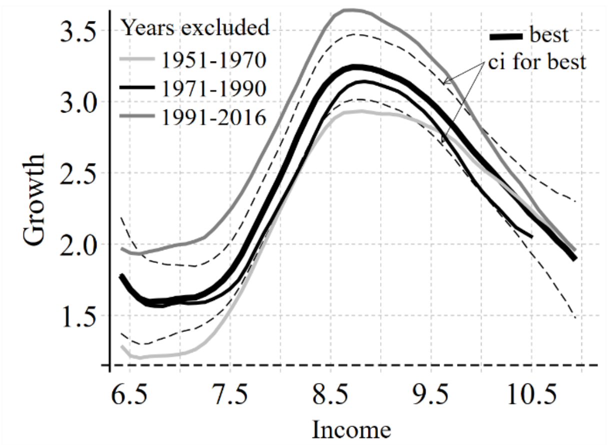

Figure 3d. Excluding country groups

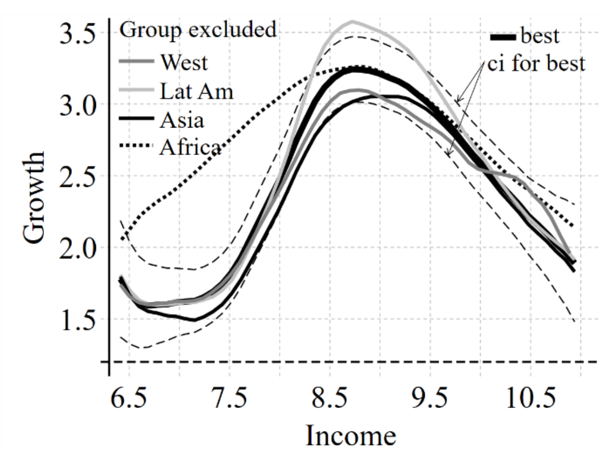

Figure 3e. Main vs OPEC

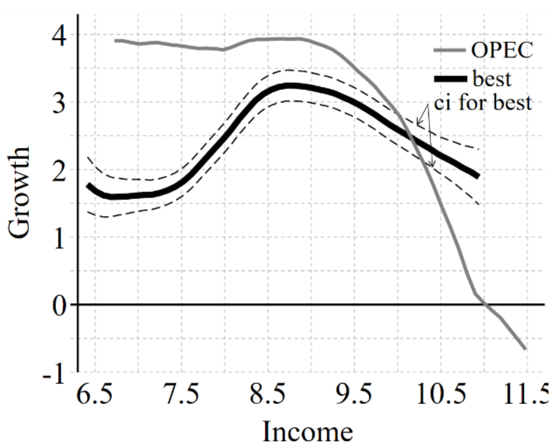

Figure 3f. Regime types

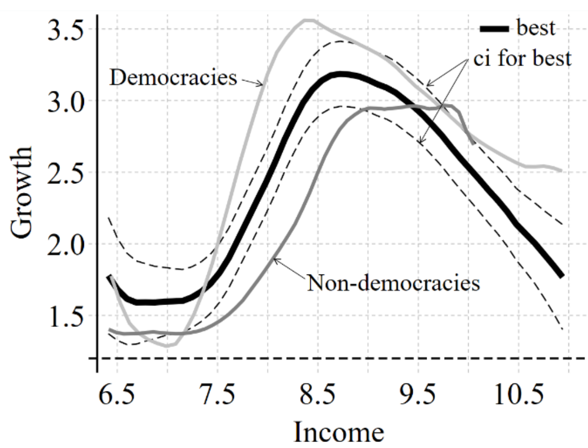

Figures 3c and d report experiments showing what happens when time periods or country groups are excluded. On Figure 3c, the peak is a bit higher when the years since 1991 are excluded. On Figure 3d, the largest effect of exclusion happens when Africa (Sub Saharan) is excluded. This makes the sample quite thin at the low end, and the form changes somewhat. However, it is still hump-shaped.

Most chapters have found that the OPEC countries have a different transition. Thus, Figure 3e compares the reference curve in the Main sample with the curve for the OPEC-countries of N = 794 observations, which excludes extreme growth rates and income levels higher than 11. A few OPEC members (especially Qatar) have even had income levels beyond 11.5 for some years. With income above 11.5, the confidence interval of the estimated OPEC kernel line explodes, so the income range has been limited at the high end. Apart from the much wider confidence interval (not shown), the estimated OPEC kernel line differs from the main empirical result. Instead of a hump, the OPEC growth-income path tends to fall throughout, which is a pattern consistent with a return to the steady state after a shock (like finding oil).

Finally, Figure 3f compares transition paths for alternative political regimes. To distinguish between democracies and non-democracies, the dichotomous measure coded by Cheibub (2010) has been used. The transition graphs for the two regime types have the same hump-shape. The graph suggests that democracies grow a little faster than the non-democracies.

The growth-income path for democracies does not differ much from the reference confidence interval. Since most high-income countries (apart from some OPEC members) are democracies, the confidence interval of the kernel line for non-democracies becomes rather wide beyond income levels of about 9, and extremely wide after income levels of about 10 (widening confidence interval not shown). So the reported kernel line for non-democracies is based on a sample that excludes two relatively rich non-OPEC oil countries (Bahrain and Oman) and Singapore, which is the only non-oil high-income country that is not a democracy.

12.4 Some consequences of the hump-shaped growth diagram

All transition curves in the book have a large variation around the central kernel curve. This is also the case for the scatter on Figure 1, which has Figure 2 as the underlying transition path. However, the data sample is large, and the transition path is well determined. The fact that the central curve looks as Figure 2 has a number of consequences that are far from trivial.[2]

[2] The presentation in section 12.4 follows standard growth theory as e.g. covered in Jones and Vollrath (2013). Barro and Sala-i-Martin (1995, 2004) have a much more detailed discussion of convergence, but nothing about divergence.

One consequence deals with the concept of β-convergence in the literature on growth empirics. It starts from a simple log-linear estimation equation:

(1) $g_{it} = \alpha + \beta y_{it} + u_{it}$, the absolute convergence equation, where g is growth and y is income (ln *gdp*), while u is the residual term

When (1) is estimated on a wide cross-country sample, the sign on β tells if the countries converge or diverge. The sign is minus for convergence, and plus for divergence. Equation (1) is known as the *workhorse* model, which can be derived from a one-sector Solow-model. The model has the modern steady state as the only equilibrium. The further below the steady state equilibrium countries are, the faster should growth be. The growth is even a hyperbolic function of the deviation from the steady state. Poor countries are far below the said steady state, and thus they should grow particularly fast. Thus, the theory predicts that the sign is negative on β. However, the standard result from estimates of (1) is that β is positive, but insignificant. Thus, the countries have a vague divergence. This also follows from the fact that the income differences between countries have been growing.

If the growth diagram looks as Figure 2, the estimate of (1) will give an insignificantly positive slope precisely as found. If the sample is thin for high-end countries, the slope may be positive, and if it is thin for low-end countries, the slope may be negative, precisely as happens on Figures 11.3a and b in the next chapter. This sign pattern also occurs in studies of convergence/divergence for the provinces/states within rich and poor countries. They converge in rich countries such as the USA, and diverge in poor countries such as India and China.

Figure 4. The growth paths of China and India over the last 86 years

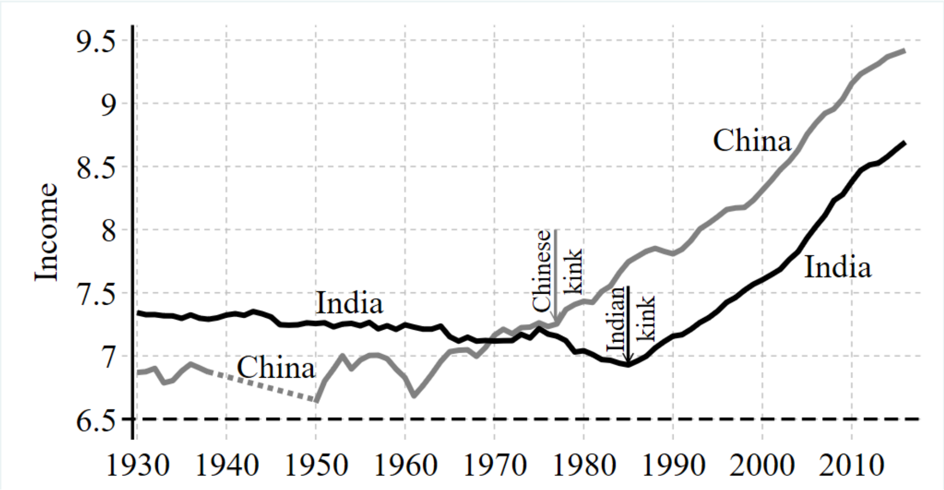

This argues that the underlying model does not apply. Solow (1965, 1970) stressed that his model was developed to explain why developed countries grow along log-linear growth paths, and return to their path even after large shocks such as wars, earthquakes or economic crises. The model is not relevant to countries in the middle of the transition, which are far from a steady state. The well-known growth paths of China and India shown on Figure 4 illustrate this. The two countries are trying to catch up. After some rather unsuccessful experiments, they both found a (similar) path that seems to work. It is not a steady state path, so growth cannot be understood as cycles around a steady state path. In addition, both countries are undergoing massive structural change.

As (1) did not show convergence, it was amended into equation (2) that contains the []-set of variables that is chosen to control for country heterogeneity. Alternatively, the constant, α, can be broken into fixed effects for countries.

(2) $g_{it} = \alpha_{it} + \beta y_{it} + [\gamma_1 z_{1it} + \ldots + \gamma_n z_{nit}] + \varepsilon_{it},$ the conditional convergence equation

With the right control-set or fixed effects for countries, β does become significantly negative. This suggests that if countries were alike, they would converge, which seems to be a tautology. However, since they are not, they do not converge. This opens up a long story, which will not be told at present. It concentrates on the non-linear form of the transition curve.

12.5. Two-sector models generate hump-shaped development paths [3]

Several chapters in the book have claimed that two-sector models where a traditional *T*-sector and a modern *M*-sector co-exist give a picture with a growth peak in the middle. Figure 5 is a sketch of the basic model.

If there are no flows between the sectors, the growth rate in the economy is just the weighted sum of the two growth rates. As the growth rate in the modern sector is the same as the one of the fully modernized rich countries, and the growth of the traditional sector is (much) lower, this ‘internal’ growth rate is lower than the growth rate of the rich countries.

The flows between the sectors open the possibility for extra growth that is potentially rather large. As explained in Chapter 1.7, a transfer of 1% of the labor force from hidden unemployment in the traditional sector to employment in the modern sector generates a growth premium of 5-8%, if the productivity gap is 5-8 times. Thus, the growth rate increases from

[3] The two-sector (or dual) model of development goes back to Lewis (1954) and Ranis and Fei (1961). For a lucid survey of the origin and later development of the model, see Gollin (2014).

below 2% to about 8%. Normally, the sectoral transfer is smaller, and many problems may occur; see Chapter 13.

Figure 5. The basic two-sector (dual) model, with dual labor markets

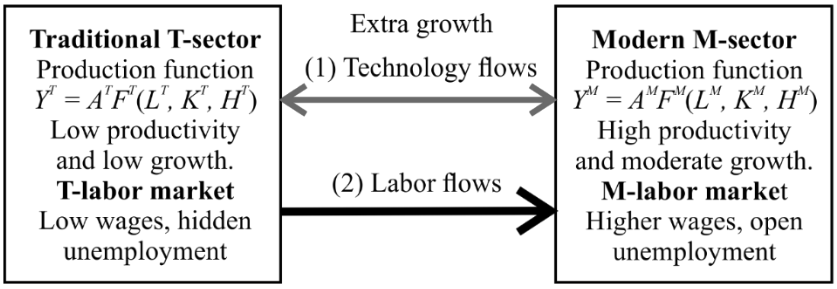

Recall from Chapter 11 that *Y* is GDP, *A* is knowledge, *L* is labor, *K* is real and *H* is human capital. The number of equations grows rapidly once we leave the one-sector approach, e.g. by adding the relations to the rest of the world, capital flows, or policies that are (much) easier to implement in the *M*-sector. The original Solow model disregarded *H*, and made *L* and *A* exogenous, so all that was necessary was the production function and an accumulation function for *K*, for everything to be solvable. With two sectors, this doubles and doubles once again when flows between the sectors are added. In addition, the two-sector model with flows makes it difficult to set wages equal to marginal productivities in the two sectors.

It is also a political temptation to use the potential growth to generate rents to distribute to clients building support for the regime. With enough rent-seeking, the growth premium vanishes; see Krueger (1974, 1990). The main theoretical property of the two-sector model is that it has no steady state as long as both sectors exist. It starts from the traditional steady state before the modern sector appears,[4] and it ends in the modern steady state when the traditional sector is fully absorbed. In order to exist as a disequilibrium, it needs some brake that limits the flows between the sectors.

(1) The technology flows go both ways. With large differences between the prices of labor and capital in the LDCs and DCs, the optimal choice of techniques in the *M*-sector differs in the two types of countries, so bits and pieces of traditional technique will be used in the modern sector, and some modern techniques will seep into the traditional sector.

(2) The resource flows, where labor is the main one, go from the traditional to the modern sector. Obviously, the possible wage gain from the sectoral migration is a driver of the labor flow. However, it is offset by the welfare losses from moving from hidden unemployment in the traditional sector to open unemployment in the modern sector, as modelled by Harris-Todaro (1970) that gives temporary equilibrium solutions for the flows.

[4] An important point to note is that the *M*-sector produces by international techniques imported from abroad.

The traditional sector has low wages, and even when wages in the modern sector are higher, they are (much) lower than in the high-income countries abroad. Thus, the modern sector has a competitive advantage at the world market. If it can sell its product, it will be rather profitable and can expand rapidly, absorbing resource flows from the traditional sector – notably labor. This might still be the case even when labor has less human capital and the goods produced are a bit behind in style and technological refinement.

A much-analyzed complication occurs when some of the firms in the *M*-sector are branches of (large) foreign firms from the DCs. Such companies are under political pressures to pay similar wages in the LDC branch as they pay at home. In addition, they often make product chains where the most labor-intensive part of the production is located in the low-wage LDCs. In other versions of the model, the *T*-sector is the export sector, while the *M*-sector consumes the import.[5] Thus, the two-sector model has been developed in many directions, and maybe there are more than two sectors.

At present, I want to argue that the whole family of such models normally produces a hump-shaped transition path. At the start, when the modern sector is small, it cannot absorb more than a tiny flow. In addition, it takes time to develop the market for the modern goods and the necessary human capital for mass production. At the end, when the traditional sector is small, it cannot generate much of an outflow. In addition, once resources are squeezed out from the traditional sector, it is likely that it will modernize rapidly, so that the sectoral gap vanishes.

Thus, it appears very likely that the greatest potential for a rapid modernization and high growth occurs somewhere in the middle, where both sectors have a substantial size. It is often assumed that the traditional sector is agriculture, so the slope of the curve for the Agricultural Transition as estimated by the kernel in Figure 1.1a illustrates the size of the flow. The slope is low at the start of the transition and high between y = 7.5 to 8.5, and as income increases, it falls.

12.6 A new two-sector model that does produce a hump-shaped development path

Two-sector models became increasingly complex during the 1980s, and then they disappeared from the literature during the 1990s, where interests changed to endogenous technical progress, processes of substitution between capital, labor and resources, etc. In order to handle these issues, the complexities of the two-sector models were unnecessary. Thus, the literature turned back to the workhorse one-sector Solow model. In the process, many authors seem to forget that the

[5] One version of the distinction is to speak about the *informal* and the *formal* sector, where firms in the informal sector have no legal title to their business. This makes it difficult to use property as collateral for loans from the banking system. This acts as a barrier for the growth of businesses.

Solow model was not made to explain development.

The two-sector model made a brief reappearance with Lucas (2009). The Lucas two-sector model contains a number of innovations relative to old two-sector models. It is made so that it is easy to collapse the model to the one-sector workhorse model. The model has “city” and “farm” as the two sectors, which both produce a single output good that adds up to GDP. Cities are the centers of intellectual exchange. The contribution of the city sector to GDP depends on the level of human capital multiplied with its employment share; it is assumed to generate a positive *agglomeration externality* due to the exchange of productive ideas in cities.

In addition, the city sector produces a *productivity externality* that spills over to the farm sector. This makes the farm output and its employment share functions of the level of human capital in cities. Assuming mobility of labor across sectors, the model predicts a declining share of farm employment with rising levels of human capital.

Growth enters the two-sector model in the form of catching up with a frontier economy, which is assumed to grow at a constant rate. As in the workhorse model, the income distance to the frontier has a positive effect on the growth rate of the follower economy, but this effect is assumed to be conditioned by an *openness externality*, such that more open follower economies should grow faster than more closed follower economies, all else constant. The model is sufficiently parsimonious as to parameters to allow a complete set of simulations that can be reported within a paper; see Gundlach and Paldam (2020). The simulations demonstrate that *all* non-collapsed versions of the model have a hump-shaped growth path.

12.7 Non-linear panel regressions with fixed effects

The kernel regression that generated the main empirical result above omitted many potentially important variables and the panel structure of the data. To address both concerns, the hump-shaped growth path is approximated with panel regressions that include country- and time-fixed effects together with a quadratic income term to allow for non-linear effects. In addition, it is checked if the marginal income effects change from positive to negative with rising levels of income, as predicted by the kernel regressions. In all regressions, the statistically significant coefficients are positive to income and negative to squared income. As expected, the explanatory power is low.

Column (1) of Table 2 gives the results for Pooled OLS, which serve as a point of reference. The marginal effects are calculated at income levels that can be directly compared with the income levels in Figure 2. The marginal income effects change as predicted by the kernel regression: positive at low-income levels and negative at high-income levels, and larger in absolute value at both ends (at y = 7.3 and 10.1) than near the peak of the hump (between 8.3 and 9.2). The negative coefficient of 0.008 at the high-income end implies a rate of convergence of about 1%.[6]

Table 2. Non-linear panel regressions

	Dependent variable: annual growth rate in %			
	(1)	(2)	(3)	(4)
Income	**6.63**	**8.90**	**7.55**	**9.59**
	(0.9)	(1.3)	(1.9)	(1.9)
Income squared	**-0.37**	**-0.51**	**-0.47**	**-0.65**
	(0.1)	(0.1)	(0.1)	(0.1)
Observations	9137	9137	9137	9137
Countries	153	153	153	153
R-squared (adjusted/overall)	0.01	0.07	0.00	0.02
Country-fixed effects	no	No	yes	Yes
Time-fixed effects	no	Yes	no	Yes
Marginal income effects at:				
y = 7.3 ($1,500)	**1.25**	**1.46**	0.63	0.15
	(0.1)	(0.2)	(0.4)	(0.4)
y = 8.3 ($4,000)	**0.53**	**0.46**	-0.30	**-1.12**
	(0.1)	(0.1)	(0.2)	(0.2)
y = 9.2 ($10,000)	-0.15	**-0.47**	**-1.17**	**-2.30**
	(0.1)	(0.1)	(0.2)	(0.3)
y = 10.1 ($25,000)	**-0.82**	**-1.40**	**-2.04**	**-3.48**
	(0.2)	(0.3)	(0.3)	(0.4)

Cross-country panel data, Main sample. Regression constant not reported, robust standard errors in parentheses.

Column (2) reports results for the inclusion of time-fixed effects, which eliminates from the sample the effects of common shocks but retains the cross-country variation. Like Pooled OLS, this specification produces a reasonable approximation of the growth path identified by the kernel regression: the marginal effects are estimated to be statistically significantly different from zero and have the right signs and relative sizes for both sides of the hump. The implied convergence rate at the high-income level is about 1.5%, but not much larger than the implicit divergence rate at the low-income level, which implies a net convergence rate close to zero.

The results change with the introduction of country-fixed effects in column (3). Eliminating the cross-country variation from the sample is like assuming that all countries are the same except for their income level, so it is almost by default that the statistically significant

6. The convergence rate, λ, can be calculated from the estimated regression coefficient (b) as $\lambda = -\ln(1+b)/t$, with t = 1 for annual growth rates.

marginal income effects are all estimated to be negative. At the high-income level, the negative coefficient of -0.02 implies a convergence rate of about 2%, which is in line with results reported in the conditional convergence literature noted above. Column (4) reports results for the inclusion of both country- and time-fixed effects. Not surprisingly, the marginal effects are much like the marginal effects estimated with country-fixed effects only.

Taken together, the results in Table 2 confirm the hump-shaped growth path of Figure 2 if the cross-country variation is maintained (columns (1) and (2)). In addition, they confirm the results of the conditional convergence literature if it is eliminated (columns (3) and (4)). Not controlling for obvious cross-country differences, except for the level of income, as in the first two specifications, will necessarily produce an omitted variables bias. But eliminating all cross-country variation, as in the latter two specifications, may be too much of a good thing, especially when assessing a potential pattern of long-run growth and development. After all, the long-run information appears to be in the cross-country variation of income levels, not in within-country variation of growth rates over time.[7]

The Grand Transition view relies on both cross-country and on time series evidence. Treating the cross-country variation as a source of omitted variables bias must lead to a rejection of the grand transition hypothesis for the sample at hand, because the within variation of growth rates in 1950-2010 does not suffice to capture the transition from a static to a modern steady state for individual countries. Maintaining the cross-country variation helps to identify a hump-shaped transition path both with kernel and panel regressions. The level of income only explains a tiny fraction of the observed variation in growth rates across countries and over time, but ignoring the Grand Transition pattern means missing a signal in the noise.

12.8 Conclusion

Kernel regressions based on cross-country panel data reveal a hump-shaped transition path for the growth rate. This empirical result contrasts with the prediction of a hyperbolic growth-income path derived from the workhorse model of growth empirics. The kernel regression results suggest that understanding long-run development calls for a two-sector model that can generate a hump-shaped growth-income path.

7. Hall and Jones (1999) use this argument to motivate their cross-country regressions on the effect of institutions on long-run economic performance. Along the same lines, Frankel and Romer (1999) use cross-country regressions in *levels* to estimate the effect of trade on (long-run) growth. The combination of persistent country characteristics and non-persistent within-country growth rates, which has been emphasized by Easterly *et al*. (1993), also speaks against eliminating *all* cross-country variation from the sample, because otherwise nothing but regression to the mean may be left.

The simulation results referred to in section 12.6 show that the hump-shaped path can be generated with a rather broad range of parameters and initial conditions, which determine the timing and the size of the hump. With obvious variation in initial conditions and possible variation of parameters across countries and over time, it becomes understandable why it has been difficult to identify a common pattern of long-run growth and development, especially with a model that imposes the restriction of a hyperbolic growth-income path. The kernel regressions reveal a common pattern of income dynamics that is overlaid by otherwise extremely noisy data; i.e., most of the enormous variation of observed growth rates remains unexplained.

The main empirical result is that the growth-income path is hump-shaped. This is supported by a number of robustness tests. If the long-run path of income can be considered as a transition from a traditional to a modern steady state, it follows by implication that the corresponding growth path must be hump-shaped. Such a growth path can be simulated based on a model that includes a traditional and a modern sector. Taken together, the hump-shaped growth-income path can be taken as the general pattern of long-run development.