The Next Hundred Years of Growth: Growth and Convergence

Richard Startz*

Revised September 2016

Abstract

I use a Bayesian Markov-switching model to forecast world GDP per capita over the next 100 years. The switching model estimates the probability that a country is currently on a path to converge to the world frontier for each country in the world, as well as the probability that a country that is not currently converging will switch to a convergent path. Forecasts depend on both the rate of growth in income in countries at the world frontier and the rate at which other countries converge to that frontier. Switching to convergence is a major source of growth. I forecast world income per capita to grow over the next hundred years at an annual rate of 2.5 percent. World income will be triple today's U.S. income, although there will continue to be considerable inter-country inequality.

JEL codes: D04, C11

Preliminary: Please Do Not Distribute

---

* Department of Economics, University of California, Santa Barbara, startz@ucsb.edu. I am grateful to Jeremy Piger for the code used in the multi-move Bayesian algorithm discussed below.

How well-off will the world be in 100 years, as measured by per capita GDP? The answer depends on growth in the world productivity frontier and on the rate at which poorer countries move toward that frontier. The latter is key: Which countries are already on a convergent path; What is the probability of countries switching to the convergent path; and What is the rate of convergence once on that path? To motivate the empirical work, I begin with a variant on the theoretical model in Lucas' (2000) "Some Macroeconomics for the 21st Century." The central features are that all countries are predicted to converge eventually, but that at a point in time a country may not yet be on a convergent path. In a given year there is some probability that a non-converging economy will switch onto a convergent path. After this, there is some rate at which that country moves toward the frontier.

Lucas answers these questions with a "reduced-form," calibrated model. In contrast, a large econometric literature has shown that, looking backwards, some countries appear to be converging toward the world frontier but others appear to be non-convergent, and has attempted to sort out causality for determining which countries are converging.[1] I take an approach somewhere in the middle. I maintain Lucas' reduced-form set-up and also maintain his assumption that *eventually* countries begin to converge to the frontier. The world is far more integrated today than it was at the beginning of Lucas' time frame, 1800, or even at the beginning of the data used in much empirical work, 1950. While there may be the occasional growth disaster in the future, it seems likely that over a period as long as 100 years most

[1] See Temple (1999) for a detailed discussion and many references. See also, Jones and Vollrath (2014, section 3.2) and Acemoglu (2009, sections 1.3-1.5) and Barro and Sala-i-Martin (1995, Chapter 12).

countries will have at least started to converge. But since I estimate the probability of a being on a convergent path, the data decides whether convergence is quantitatively important or not.

My approach also picks parameters in a manner that is mid-way between Lucas' calibration and the unrestricted estimation used in much of the literature. I employ a Bayesian framework, positioned part way between calibration and purely data-driven estimates. The Bayesian framework allows one to impose assumptions that limit the range of parameters. In particular, this helps the model pick out low frequency rather than business cycle changes. The Bayesian framework forces one to be explicit about what assumptions are made and I specify priors that are informed by what we know about growth. At the same time, the priors are loose enough that the data chooses parameter estimates within a broad range.

The heart of the estimation is a Markov-switching model that estimates how likely it is that a given country is already on a path toward convergence, and uses the estimated history of countries switching from non-convergent to convergent to determine the probability that not-yet-converging countries will switch in the future. I then apply this probability, together with an estimate of the rate of convergence and the rate of frontier growth to forecast income out into the future. I estimate the model over the post-War period and forecast using the posterior distribution of the parameters of the model and the posterior distribution of the state of convergence for each country.

In order to see why estimating convergence is central to forecasting, it is useful to begin with some back-of-the-envelope calculations. Frontier income has been growing at about 1.7

percent per year since 1970.[2] Of course, that growth rate might change in the future. Gordon (2016, p. 14) gives some estimates that give us an idea of plausible ranges for growth in the frontier. He reports that U.S. income grew at 1.8 percent from 1870 to 1920 and 2.4 percent from 1920 through 1970. Notably, Gordon argues that future growth won't exceed the current rate. Forecasting frontier growth over the next 100 years is a difficult and open question, as there might be further productivity slowdowns or there might be quickened growth due to increased automation. Take 1.7 percent as a comparison number if there were to be no further convergence of countries to the world frontier.

The *potential* for faster world growth through convergence is enormous, because most of the world is still very far from the frontier. Taking the United States as representative of the world frontier, as is customary in the post-War period, frontier income is now $51,600 while rest of world income is $9,800. If the rest of the world were to catch up to the frontier in a single, 20-year generation (miraculously), then the world growth rate would equal 9.1 percent. If convergence were complete in 50 or 100 years respectively, then the world growth rate would have to be 465 or 3.2 percent. All these numbers are much larger than plausible rates of frontier growth.

Of course, theory does not dictate that the next century should see complete convergence. Indeed, I estimate income in the rest of the world will rise from 25 percent of the frontier to 58 percent of the frontier level over the next hundred years. This suggests an overall growth rate

[2] Data is from 1950 through 2014 in 2011 dollars, as given in the Penn World Tables 9.0 including updates through August 2016, 2016. See Feenstra et. al. (2015). "Frontier income" means U.S. GDP per capita throughout, while "rest of world" means all countries other than the U.S.—including other countries with income that is also at the frontier.

of 2.5 percent, which is about half again speedier than frontier growth. Predictions are shown in Figure 1. Without convergence, I estimate a hundred years of growth will put world income about sixty percent higher than today's U.S. income. With estimated convergence, world income will be more than triple today's U.S. income. (U.S. income will be nearly six times today's income.) These are median point estimates; fiducial intervals are presented below.

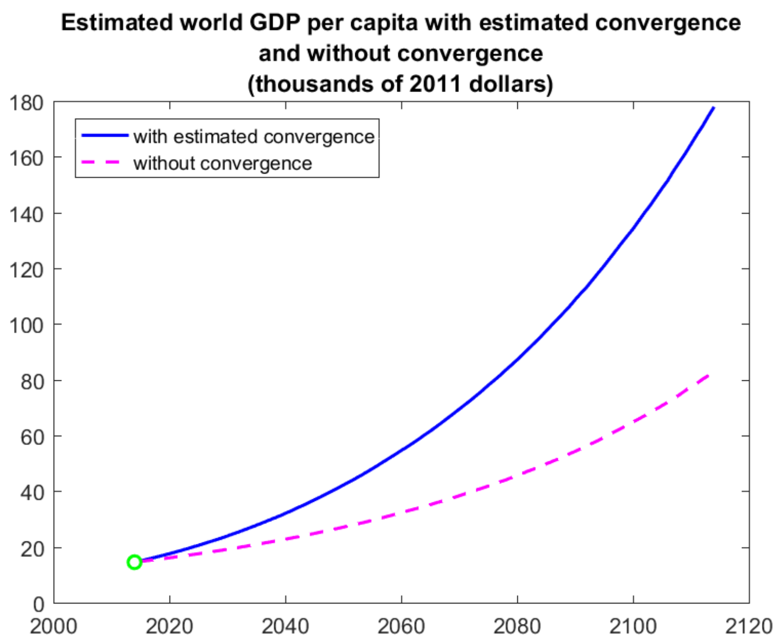

Figure 1

Turn now to the importance of using priors to help identification. In unrestricted estimation of state-switching models, it is difficult to separate out low frequency convergence from higher frequency issues related to business cycles. Figure 2 illustrates the issues for the Republic of Korea and Spain.

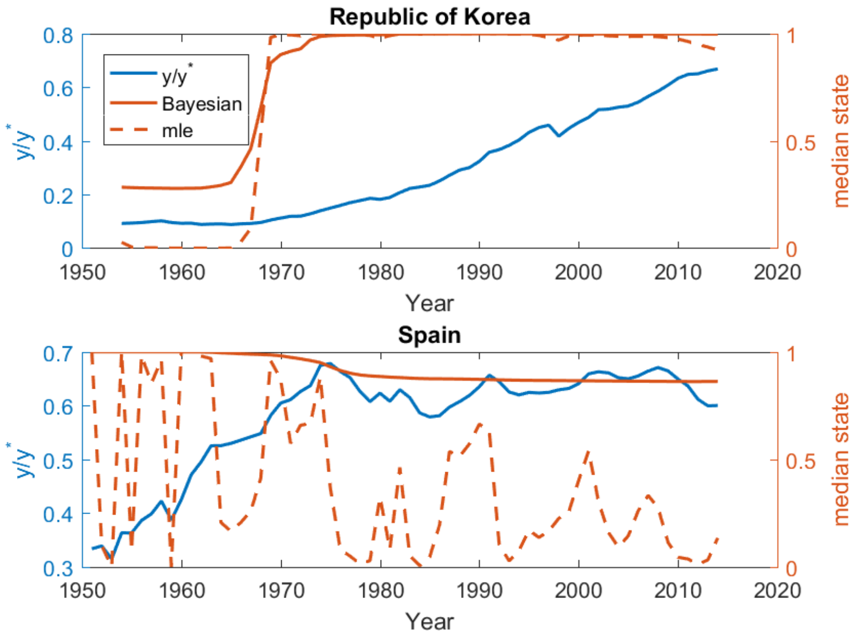

Figure 2

The upper panel in Figure 2 shows the ratio of income in Korea as compared to the frontier. At some point, perhaps in the late 1960's as one of the "Asian tigers, Korea began to slowly converge toward frontier income. The line marked "mle" shows the maximum likelihood estimate of the (smoothed) probability that Korea was on a converging path, as estimated from a standard Markov-switching, maximum-likelihood approach. A visual comparison of the changing slope of the ratio-to-frontier and the estimated probability of being on a convergent path suggest the two are reasonably consistent. The line showing a Bayesian estimate of the probability of convergence sends essentially the same message.

The lower panel in Figure 2 gives the analogous estimates for Spain. Through some time in the 1970s Spain appeared to be converging with the frontier, after which growth leveled off and convergence may have stopped. Unlike in the upper panel, the maximum likelihood estimates of convergence fluctuate wildly. The problem is that the low frequency component is not being well-identified. In contrast, the Bayesian estimates are well-behaved. In the early period, Spain seems to be on a convergent path. Later the probability falls, although continuing to indicate that convergence is more likely than not.

The difference between the mle versus Bayesian estimates comes from the fact that the Bayesian estimates include identifying restrictions that help pick out low frequency behavior. In other words, there is nothing important in the use Bayesian techniques per se, it is simply much easier to impose the needed identifying restrictions. While discussed in more detail below, the essence is that I specify priors in which switches between convergent and non-convergent states are relatively rare, and in which the rate of convergence while in a convergent state has to take on a “reasonable” value.

In the next section, I present an illustrative model to nail down concepts and compare the model to Lucas’ (2000). I then discuss the data and the estimation technique, including the role of the identifying restrictions. I then make the kind of estimates shown in Figure 2 for all the countries in the world and use the estimates to project income for each country for the next century.

I. Models and related literature

Projections of long-run world growth are often done using “scenarios” or calibration, unlike the purely econometric methods used here. And unlike the techniques here forecasters often project out factor inputs and then run these inputs through a production function. Examples include Johansson et. al. (2012) and Dellink et. al. (forthcoming). The econometric estimate closest to the one here is Zimmer et. al. (2016), which estimates convergence rates for each country in the world based on historical convergence rates. Using Bayesian techniques somewhat similar to those used here, but assuming that countries never change convergence paths, Zimmer et. al. (2016) estimate that for over half the countries the rate of convergence is zero, so world income is projected to grow at little more than the rate of frontier growth.

The paper with the closest model to the one used here is Lucas (2000). Lucas offers a calibrated model in which the frontier grows at a specified rate. Lucas, focusing on the period since 1800, has countries either in no growth status or in convergent status. A given country has no growth until it begins convergence, which happens probabilistically. The estimates here, based on the modern era, look at growth at the frontier rate without or with convergence. The mechanics are the same when there is convergence, but the alternative is no growth in Lucas or conditional convergence here.

I present here a simple model to illustrate the assumed mechanics. Let y be log per capita GDP. A country has two production sectors. Each country has one unit of resource per capita, divided between the developing sector and the modern sector, $X_D + X_M = 1.0 \leq X_D, X_M \leq 1$. Both sectors are exposed to the world technology frontier, A, but the developing sector makes inefficient use of the technology. Production is $Y_D = \beta A X_D, 0 < \beta < 1$ and $Y_M = A X_M$. For countries at the world technology frontier we have $X_M = 1$,and income $Y^* = A$. For other countries, $Y = Y_D + Y_M = \beta A X_D + A X_m$. The ratio of a country's income compared to the world frontier is $Y/Y^* = \beta X_D + X_M$, so the growth rate of the ratio is

$$[Y\dot{/}Y^*] = (1-\beta)\dot{X}_M \tag{1}$$

There is a continuum of agents endowed with time that can be turned into effortlessly into one unit of X_D in aggregate. Alternatively, agents can engage in costly effort to create X_M. The cost of engaging in the modern sector is decreasing in the size of the existing modern sector. The dollar-equivalent utility cost of creating modern input for agent j is $\lambda_j - \ell X_M$, where λ is distributed across the population with density $f(\lambda)$. An agent is indifferent between engaging in the two sectors if $\beta A = A - (\lambda_j - \ell X_M)$, so the marginal agent has cost

$$\lambda^* = A(1-\beta) + \ell X_M \quad (2)$$

Allocation of inputs is determined by

$$X_M = \int_{-\infty}^{\lambda^*} f(\lambda) d\lambda \quad (3)$$

If the cost of the marginal agent is below the lower support of $f(\lambda)$, $X_M = 0$ and the economy does not converge. If λ^* is greater than the upper support $f(\lambda)$, $X_M = 1$ and the economy is on the frontier. In either case, the gap between country income and frontier income is constant. In between, the gap closes according to equation (1). If one assumes for convenience that $f(\lambda)$ is uniform with density c inside the support, then from equation (3) we have $\dot{X}_M = c\dot{\lambda}^*$. From equation (2) we have $\dot{\lambda}^* = \dot{A}(1-\beta) + \ell \dot{X}_M$. Together with equation (1) we have that for interior solutions the gap closes at rate

$$[Y \dot{/} Y^*] = (1-\beta)\frac{c(1-\beta)}{1-c\ell}\dot{A} \quad (4)$$

I close the model by assuming that log frontier income, y^*, follows a random walk with growth rate $\mu > 0$. Positive frontier growth bakes into the model the result that economies eventually start to converge. In order to bring the model to the data, I simplify by assuming that all countries have the same convergence rate when they have an interior solution; linearize equation (4) and move it into discrete time; and add the possibility of random shocks. The model to be estimated becomes

$$\log\left(\frac{y_{i,t}}{y_t^*}\right) = (1 - S_{i,t}) \times \log\left(\frac{y_{i,t-1}}{y_{t-1}^*}\right) + S_{i,t} \times \rho \times \log\left(\frac{y_{i,t-1}}{y_{t-1}^*}\right) + \varepsilon_{i,t} \quad (5)$$

$$y_t^* = \mu + y_{t-1}^* + \varepsilon_t^* \quad (6)$$

where countries are indexed by $i = 1, \dots, n$. $S_{i,t} \in \{0,1\}$ equals 1 during the convergence period, so the log ratio between a country's income and the frontier closes at an expected rate of μ percent. Outside the convergence period, $S_{i,t} = 0$, the gap is unchanged expect for random fluctuations. Note that equations (5) and (6) are very close to equations (1) and (2) in Lucas (2000).

To complete the empirical specification, I take $S_{i,t}$ to follow a discrete first-order Markov process with the transition probabilities p_{00} (to be estimated) and $p_{11} \approx 1$. (Jones (1997) also uses a Markov specification in a somewhat different way to study the distribution of world income. Jones emphasizes the importance of convergence in determining the world income distribution. Durlauf and Johnson (1995) divide countries into different growth regimes, although not using a Markov-switching model.) Notice that in this specification that a country which has caught up to the frontier is expected to stay there in both states, except for random fluctuations.

II. Data

The data used here is from the Penn World Tables (Feenstra et. al. (2015)) version 9.0. GDP per capita is measured in 2011, PPP converted, U.S. dollars and is measured according to the PWT variables $Y = RGDPNA/POP$. The data measures income for 181 countries annually with various starting dates beginning in 1950 and all ending in 2014. Table 1 provides descriptive statistics of the underlying data, showing both population-weighted and unweighted descriptive statistics.[3] Observations start in a variety of years; the weighted mean first observation comes in 1955. The shortest observation period is 10 years; the longest is 65. The ratio of average income to frontier income declined in the data set from individual country starting dates through 2014, but this largely reflects later entry into the data set of poorer, smaller countries. Note that the population-weighted ratio of between country income to the frontier closed modestly between the first-observation for a country and 2014. The high maximum income ratios are due to Qatar and Macao.

	Mean (Weighted by 2011 GDP)	Mean (unweighted)	Min	Max
Starting year	1955	1963	1950	2005
Average growth rate, rest of world	2.70	1.88	-2.67	7.14
Average growth rate, U.S.	1.99			
2014 income, rest of world	$12,747	$19,132	$570	$163,294
2014 income, USA	$51,621			
Ratio of income to US income, first observed year	0.14	0.68	0.01	16.5
Ratio of income to US income, 2014	0.25	0.37	0.01	3.16
Source: Penn World Tables 9.0. Data in 2011, PPP-adjusted U.S. dollars				

Table 1

PWT provides an unbalanced panel. There is likely some selection bias in terms of when coverage begins for a particular country. However, 89 percent of the world's population, by 2014 population weights, is covered by 1960. Figure 3 shows population coverage as well as the population-weighted, observed gap from the frontier, $y_{i,t}/y_t^*$. The average country experienced a slowly closing gap except for the first two years, which covered a much smaller fraction of the population than did the remainder of the sample.

[3] Weighting is done throughout using 2014 population weights with the universe defined as the countries observed in PWT.

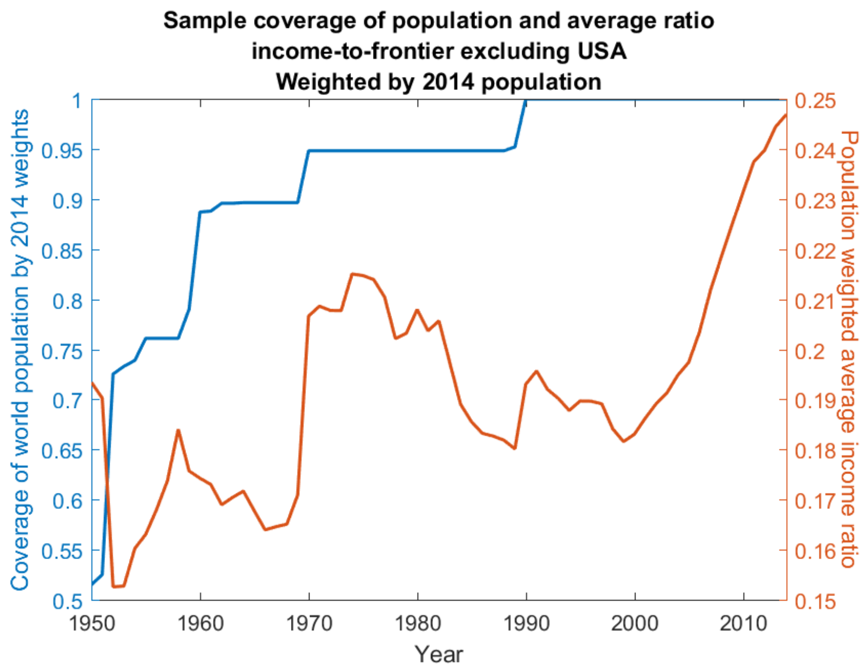

Figure 3

One question is whether the data is broadly consistent with a model of all countries having positive growth, except for random fluctuations. Mean growth was negative for 22 countries, although significantly negative with a one-tailed five percent test for only three. As a practical matter, the countries with negative growth are very small, accounting for less than four percent of world GDP in 2014. The model also expects a country's growth to be equal to or higher than frontier growth, again with the exception of random shocks. Quite a few countries, 95, had mean growth lower than the frontier, but only 26 had growth significantly below the frontier on a one-tailed, five percent test of the difference between country and U.S. mean growth rates. Total 2014 GDP of the countries with growth significantly below the mean was 2.4 percent of world GDP, with a quarter of that accounted for by South Africa.

III. Estimation

Estimation follows from a fairly standard Bayesian algorithm, which is described first. I then describe the important identifying restrictions. The model consists of 180 equations for the income gap plus an independent random walk equation for frontier income, as specified in equations (5) and (6).

The frontier equation is estimated assuming normal errors and independent normal-gamma priors.[4] The priors are diffuse, so that the posteriors mimic the standard frequentist distribution for a sample mean. μ and $h_f \equiv 1/\sigma^2_{\varepsilon^*}$ are estimated by Gibbs sampling with 100,000 samples retained after discarding 10,000 samples. The significant issue is that the U.S. underwent a growth slowdown somewhere around 1970; Gordon (2016) uses the date 1970 and Perron (1989) prefers 1973. Since projections for the future should reflect current growth, I estimate μ beginning in 1970. Prior and posterior statistics are given in Table 2. The posterior mean for μ matches the sample mean growth to two digits.

Parameter	Conditional prior	Posterior mean	Posterior standard deviation
μ	$N(0,.0018) \times \Gamma((4.6 \times 10^{-4})^{-1}, 1)$	0.017	0.0031
h_f		2424	510
ρ	$U(0.95,0.98)$	0.97	0.0093
$\bar{h}$	$\Gamma(48,0.125)$	481	1014
p_{00}	$\beta e(39,1)$	0.997	0.0013

Table 2

[4] There are several parameterizations of the gamma distribution in the literature. I use $f_\Gamma(z|m,\nu) \propto z^{\frac{\nu-2}{\nu}} \exp\left(-\frac{z\nu}{2m}\right)$.

The country equations in equation (5) are estimated by a Gibbs sampler, again with 100,000 samples retained after discarding 10,000 samples, iterating over conditional draws for the state vector S_i for each country, the precision $h_i \equiv 1/\sigma_{\varepsilon_i}^2$ for each country, and common values for ρ, and p_{00}. Holding the last two parameters in common permits identification for countries in which most estimated states are convergent or most estimated states non-convergent. The Gibbs sampler is standard, with the conditional state draws made by a multi-move algorithm. (See Carter and Kohn (1994) and Kim and Nelson (1999).)

Conditional on S, I assume independent normal-gamma priors for ρ and h_i. The priors for the precision are diffuse, with the prior mean for each country set to the unconditional variance of $\log\left(y_{i,t}/y_t^*\right)$ and the prior standard deviations set to four times the prior means. The across-country average values of the priors and posteriors are given in Table 2.

The advantage of the Bayesian technique is that it allows for imposing identifying restrictions and for being explicit about what is being imposed. (Gibbs sampling also allows for considerable computational simplicity.) We want state estimates that are consistent with the economic idea that at a point in time economies either are or are not on a convergent path, and that once on the path generally stay there. Avoiding picking up higher-frequency business cycle coordination requires keeping $\rho < 1$ and is aided by keeping ρ relatively high. I do this with an informative uniform distribution that bounds ρ away from 1. At the lower bound, a country would close half the income gap from the frontier (starting at the average gap) in an implausibly fast 22 years. At the upper bound, the half-life of the average gap is 55 years. While the time series are relatively short for some countries, there are many total observations. This allows for identification of the common parameters. The posterior is reasonably well-identified within the prior interval with the posterior mean 2.13 standard deviations above the lower limit of the prior and 1.09 standard deviations below the upper limit. The half-life to close the average gap at the posterior mean is 35 years.

The first-order Markov process for the state-switching model is controlled by the probabilities of remaining in the respective states, p_{00} and p_{11}, as well as a country's observed income gaps. In Lucas' model, once a country starts to converge it continues to converge. This is largely true in the model above, although a sufficiently large random shock could cause a state-switch to a non-convergent state. I set $p_{11} = .9999$. Presumably, switches into convergence are also fairly rare. Parameters are given for the beta distribution prior for p_{00} in Table 2, but perhaps are more easily understood as specifying a prior mean of 0.975 and prior standard deviation of 0.0244. The posterior is well-identified by the data; note that the posterior standard deviation is 1/20th of the prior standard deviation. At the posterior mean, the probability of a non-converging country having switched into the convergent state after 100 years is 29 percent. The steady-state probability for being in the convergent state is 97 percent. While the steady-state probability is quite sensitive to the value of p_{11}, convergence to the steady state is quite slow. The value of p_{11} makes little difference for a country that is already on the convergent path, e.g. Korea, and only a little difference for a country that is not already on the convergent path.

The Gibbs sampler generates a posterior distribution for the convergence state, S, for each country in each year. Draws for each country are done separately, and draws depend on the entire history of the respective country's log ratio relative to the frontier, as well as the Markov probabilities. The mean of the posterior S_{it} give the probability that country i has switched to the convergent state in year t. Figure 4 shows the probabilities averaged across countries. The average probability of being on a convergent path rose from 14 percent in 1960 to 52 percent in 2014. Based on the posterior estimates, the probability that Korea was in a convergent state in 2014 was 99 percent. The probability for China was 92 percent and the probability for India was 70 percent. But the probability that Pakistan had moved into a convergent state was under 2/10th of one percent.

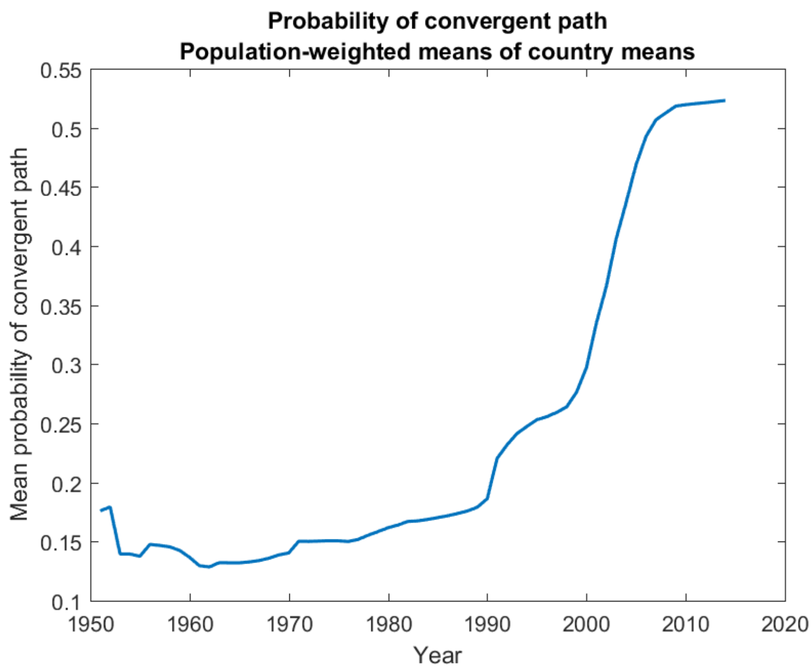

Figure 4

IV. Results

Predictions over the next 100 years are made by taking 10,000 draws from the posterior distributions and then projecting the frontier and the gap from the frontier for each country. I begin with a draw for μ and h_f. The frontier is then forecast by taking 100 normal draws for ε^* and running out the random walk with drift in equation (6). Next, draws are taken for $\rho, p_{00}, h_i,$ and $S_{i,2014}$. With respect to the latter, note that while each state draw is either zero or one, the expected value of each draw is the posterior mean for each country. The states are then forecast out according to a first-order Markov process governed by p_{00} and p_{11}. Finally, 100 normal draws are taken for $\varepsilon_{i,t}$ and income is projected for the country using equation (5).

As shown in Figure 1, the median forecast for world income in 100 years is $178,000. This compares to a forecast frontier income of $296,000. In contrast, absent further convergence world income would be predicted to rise only to $83,000. Of course, each of these numbers would be changed if the average growth rate of the frontier over the next 100 years turns out to be significantly higher or lower than the average growth rate over the last 40.

In order to better understand the dynamics over the next 100 years it is useful to separate the effects of convergence along the convergent path from the effects of switching from a non-convergent to a convergent path. Consider a country that is not currently converging. While it is there is a reasonable likelihood that it will begin to converge over the next 100 years, any switch will likely come too late to make much difference to growth. In Figure 5 the non-convergent line graphed on the right axis—note the scaling of the right axis—shows the mean path for a country that is initially non-convergent (e.g. Pakistan). For all practical purposes, the country never closes the gap with the frontier at all. Contrast this with a country which is already converging (e.g. Korea). Even though the rate of convergence is slow, after 100 years the country will have essentially converged to the frontier.

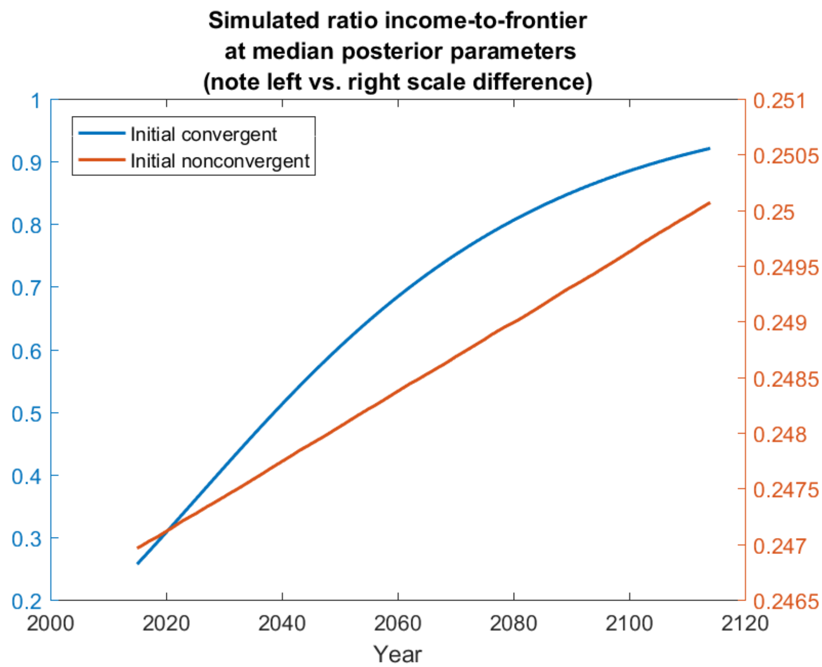

Figure 5

To a surprising extent, the importance of convergence depends on two countries: China and India. 38 percent of the difference between the forecast 100-year world income and the forecast absent convergence is accounted for by these two countries. Russia accounted for another 1.3 percent and no other country accounted for as much as 1 percent of the difference. The estimated model does not make use of idiosyncratic information about likely future growth prospects for individual countries, above and beyond the information encapsulated in the historical growth record. The fact that nearly all countries play small roles in the overall world forecast suggests mis-forecasts for individual countries are unlikely to have much effect on the final result so long as one believes that China and India are likely to continue to converge.

As seen in Figure 1—which gives the median forecast—world income is projected to grow substantially over the next 100 years and the majority of that growth comes from convergence to the frontier. Since the forecasts are based on how much countries have converged to the frontier, it is useful to look at projections for the gap from the frontier and the probability of being in a convergent state. Forecasts are shown in Figure 6. The 100-year forecast is that the gap from the frontier will have risen to 61 percent and that the probability of being in a convergent state will have risen to 65 percent. However, there considerable uncertainty about the forecasts, reflecting largely the variance of ε_{it}. At the 10th percentile of the forecast posterior the income gap will be effectively unchanged and at the 90th percentile income in the rest of the world will have overtaken forecast U.S. income.

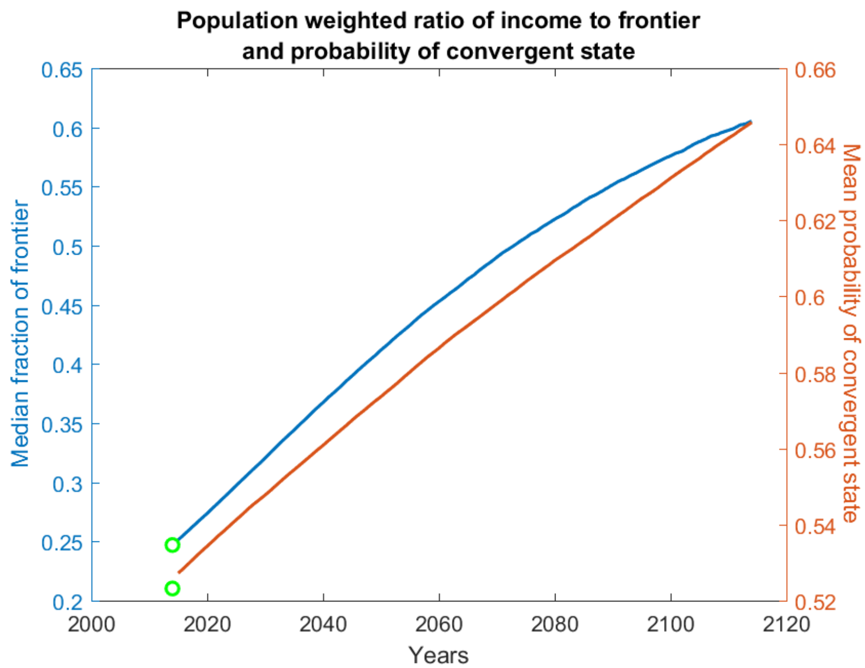

Figure 6

Uncertainty about convergence comes from two sources. The first source of uncertainty is due to uncertainty about future shocks being in the convergent state and to future additive errors in equation (5). The second source of uncertainty is due to estimation of the parameters, i.e. posterior spread. This second source can be eliminated by making projections based on median posterior values for $\{\mu, h_f, \rho, p_{00}, h_i\}$, while retaining posterior uncertainty about $S_{i,2014}$.

Figure 7 gives median projections and 90 percent highest posterior density intervals both with and without parameter uncertainty. Most uncertainty arises from real shocks and state uncertainty; parameter estimation uncertainty plays a fairly small role.

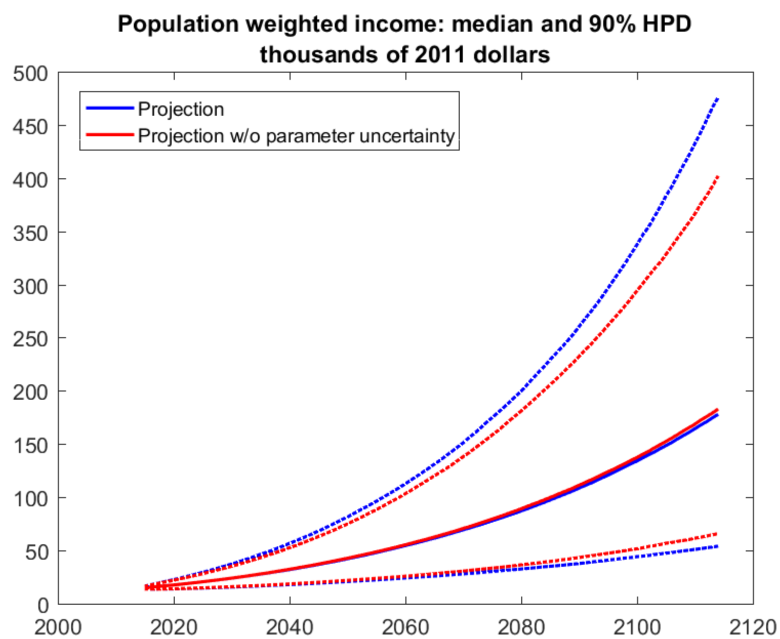

Figure 7

Lucas (2000) wrote "I think the restoration of inter-society income equality will be one of the major economic events of the century to come." While the principle topic of interest is forecasting future world income levels, the model also generates forecasts of income distribution. Figure 8 graphs the fraction of the world's population living at or below a given income level in 2014 and one hundred years later, using median country forecasts. In 100 years, most of the world's population will be in countries with far higher income levels than today's U.S. income. However, it seems unlikely that world income will have reached equality.

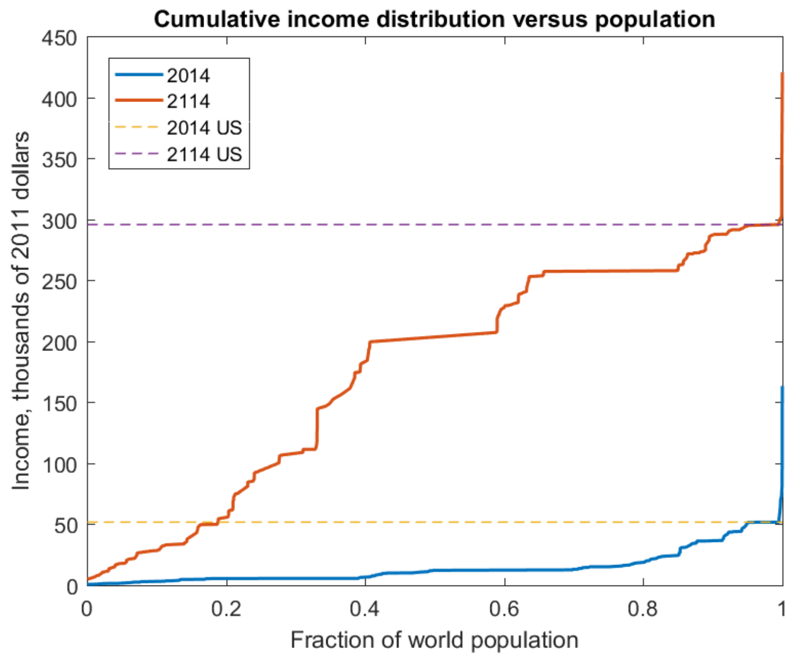

Figure 8

V. Conclusion

A good forecast for 100 years hence is that income will be roughly triple today's U.S. income and that most of the world's population will be living in countries with high incomes by today's standard. Such growth reflects a great increase in the number of countries which have begun to converge to the world income frontier, together with somewhat slow convergence once the convergence process begins. Note that this forecast is rather higher than forecasts which do not separate convergent and non-convergent states. For example, Zimmer (2016) forecasts world income growth of only 1.8 percent, in contrast to the 2.5 percent growth rate forecast here.

In 100 years, a good fraction of the world's population will live in countries with income not far below the world frontier. Some four-fifths of the population will live in countries with higher income than the United States today. The typical standard of living will be much, much improved, but inter-country inequality will remain.

References

Acemoglu, Daron. 2009. *Modern Economic Growth,* Princeton University Press.

Barro, Robert J. and Xavier Sala-i-Martin. 1995. *Economic Growth,* McGraw-Hill.

Carter, C.K. and R. Kohn. 1994. "On Gibbs Sampling for State Space Models," *Biometrika,* 81 (3): 541-553.

Dellink, Rob, Jean Chateau, Elisa Lanzi, and Bertrand Magné. Forthcoming. "Long-term economic growth projections in the Shared Socioeconomic Pathways," *Global Environmental Change.*

Durlauf Steven N. and Paul A. Johnson. 1995."Multiple Regimes and Cross-Country Growth Behavior," *Journal of Applied Econometrics,* 10 (4), 365-384.

Feenstra, Robert C., Robert Inklaar and Marcel P. Timmer. 2015. "The Next Generation of the Penn World Table," *American Economic Review*, 105 (10): 3150-3182.

Gordon, Robert J. 2016. *The Rise and Fall of American Growth,* Princeton University Press.

Johansson, Åsa, Yvan Guillemette, Fabrice Murtin, David Turner, Giuseppe Nicoletti, Christine de la Maisonneuve, Philip Bagnoli, Guillaume Bousquet, and Francesca Spinelli. 2012. "Looking to 2060: Long-term global growth prospects," *OECD Economic Policy Papers,* No. 03.

Jones, Charles I. 1997. "On the Evolution of the World Income Distribution," *Journal of Economic Perspectives,* 11 (3): 19-36.

________ and Dietrich Vollrath. 2013. *Introduction to Economic Growth, 3rd ed.,* W.W. Norton & Company.

Kim, Chang-Jin and Charles R. Nelson. 1999. *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications,* MIT Press, Cambridge, MA.

Lucas, Robert E. Jr. 2000. “Some Macroeconomics for the 21st Century,” *Journal of Economic Perspectives,* 14 (1): 159-168.

Perron, Pierre. 1989. “The great crash, the oil price shock, and the unit root hypothesis,” *Econometrica* 57, 1361–1401.

Temple, Jonathan. 1999. “The New Growth Evidence.” *Journal of Economic Literature,* 37 (1), 112-156.

Zimmer, Alec, Adrian E. Raftery, Dargan Frierson, and Richard Startz. 2016. “Less Than 2° C Warming by 2100 Unlikely,” Working paper, University of Washington.