

NBER WORKING PAPER SERIES

AN ECONOMETRIC MODEL OF INTERNATIONAL LONG-RUN GROWTH DYNAMICS

Ulrich K. Müller
James H. Stock
Mark W. Watson

Working Paper 26593
<http://www.nber.org/papers/w26593>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2019

For helpful comments we thank participants at several seminars and in particular those at the Resources for the Future workshop on Long-Run Projections of Economic Growth and Discounting. Müller acknowledges financial support from the National Science Foundation grant SES-1919336. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Ulrich K. Müller, James H. Stock, and Mark W. Watson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

An Econometric Model of International Long-run Growth Dynamics
Ulrich K. Müller, James H. Stock, and Mark W. Watson
NBER Working Paper No. 26593
December 2019
JEL No. C32,C55,O47

ABSTRACT

We develop a Bayesian latent factor model of the joint evolution of GDP per capita for 113 countries over the 118 years from 1900 to 2017. We find considerable heterogeneity in rates of convergence, including rates for some countries that are so slow that they might not converge (or diverge) in century-long samples, and evidence of “convergence clubs” of countries. The joint Bayesian structure allows us to compute a joint predictive distribution for the output paths of these countries over the next 100 years. This predictive distribution can be used for simulations requiring projections into the deep future, such as estimating the costs of climate change. The model’s pooling of information across countries results in tighter prediction intervals than are achieved using univariate information sets. Still, even using more than a century of data on many countries, the 100-year growth paths exhibit very wide uncertainty.

Ulrich K. Müller
Department of Economics
Princeton University
Princeton, NJ 08544-1013
umueller@princeton.edu

Mark W. Watson
Department of Economics
Princeton University
Princeton, NJ 08544-1013
and NBER
mwatson@princeton.edu

James H. Stock
Department of Economics
Harvard University
Littauer Center M26
Cambridge, MA 02138
and NBER
James_Stock@harvard.edu

An appendix is available at <http://www.nber.org/data-appendix/w26593>

1. Introduction

Among the many technical problems in the economics of climate change is the need to make projections of global and regional economic growth into the deep future. Future levels of GDP drive future energy consumption, future emissions of carbon dioxide, the economic capacity to reduce those emissions, and human ability to adapt to the changing climate caused by those emissions. These factors thus affect the economic cost of climate change and in particular the Social Cost of Carbon (SCC), which is the monetized net present value of the economic damages resulting from emitting an additional ton of carbon dioxide. Because of the long residence time of carbon dioxide in the atmosphere, the horizon for these projections is up to several centuries. And, because the effect of climate change is location specific, both regional and global projections are required. The method used to make the initial US Government estimate of the SCC, described in Greenstone, Kopits, and Wolverton (2013), uses five judgmental scenarios of global economic growth. It is generally recognized, however, that a modern, internally consistent probabilistic approach grounded in historical data is needed. In particular, the National Academy of Sciences (2017, Ch. 3) report on the SCC suggests using the Müller-Watson (2016) long-term trends framework to produce long-term GDP forecasts.

This paper develops a probability model of the joint growth of national per-capita GDP, estimated using up to 118 years of data on 113 countries. The premise of this exercise is that the joint stochastic process followed by the long-run growth of national incomes over the past century is a useful starting point for projecting their evolution – more precisely, for computing their joint predictive probability distribution – over the next 100 years. To this end, our model admits, but does not require, a number of features that the previous literature suggests characterize long-term economic growth. First, the model allows for convergence (in expectation) to a single global factor, however there can be considerable heterogeneity in the rate of convergence. Second, if this rate of convergence to the global factor is sufficiently slow, a century-long realization can produce apparent convergence to parallel paths in the long run (so-called conditional convergence). Third, an individual country can have a highly variable long-term growth rate, including strong multidecadal growth and prolonged periods of economic collapse. Fourth, the model allows for “convergence clubs,” that is, clusters of countries with highly correlated long-run income levels within the cluster. Fifth, the global factor itself is

allowed to evolve in a flexible way that, consistent with the historical evidence, allows for persistent changes in its underlying mean growth rate. We build these features into a multi-factor Bayesian dynamic factor model, where the factors are distinguished by their dynamics and their (latent) commonality across groups of countries.

Although we have motivated this exercise as one of long-run forecasting, our analysis also sheds new light on a number of classic questions in the large and mature literature on growth regressions and cross-country convergence (Barro (1991); for recent reviews see Jones (2016) and Johnson and Papageorgiou (2018)). In particular, we find additional evidence of historical convergence groups in the form of disconnected networks that include rapidly developing and developed countries in Asia, Anglo-Saxon countries, and selected countries in Europe, the former Soviet bloc, Africa and Latin America. As this list suggests, the model endogenously provides evidence of regional convergence. We also find considerable variation in the rate of convergence to the global factor, with a 90 percent interval of half-lives ranging from 78 years to 502 years, pooled across all countries.

Our analysis also contributes methodologically to both the low-frequency modeling and growth convergence literatures. The extension of the Müller-Watson (2008, 2016, 2018) approach from a small number of time series to the joint evolution of 113 time series is nontrivial. To make this extension we draw on both their low-frequency trend approach, Bayesian dynamic factor modeling, convergence on multiple time scales, and time series models for growth rate dynamics. By focusing on low-frequency trends, we filter out variation and comovements with periods shorter than 15 years, including comovements over business cycles.

We begin in Section 2 with a description of the data, which is an unbalanced panel of GDP per capita for 113 countries, taken from the Penn World Table (Feenstra, Inklarr and Timmer (2015)) and the updated Maddison Project Database (Bolt, Inklarr, de Jong and van Zanden (2018)) and covering 1900 to 2017. Plots and descriptive statistics highlight five features of the data: a common growth factor, persistent changes in long-term growth rates within countries, a temporally stable dispersion of the historical cross-sectional distribution, extremely persistent country-specific effects, and a possible group structure of cross-country correlations.

Section 3 outlines an econometric model that captures these features. The model has a simple structure, but it allows cross-country heterogeneity and a flexible pattern of dynamic

covariability across the 113 countries. Unfortunately, this flexibility comes at the expense of introducing hundreds of unknown parameters.

Section 4 takes up the problem of estimating these parameters and computing the long-run joint predictive distribution for the 113 countries. We focus on 50- and 100-year ahead predictions. As a general matter, Bayes estimation of a high dimensional model ($n = 113$ countries, $T = 118$ years, with over 800 parameters, with a goal of estimating a joint predictive distribution 100 years into the future) presents daunting computational challenges.

As we show, however, the structure of the model, priors, and data transformations yield important simplifications. Because the long-run nature of our analysis allows us to focus on low-frequency averages of the raw data, the effective dimension of the data is reduced by a factor of approximately seven. Furthermore, because those low-frequency averages follow normal laws in large-samples, estimation can be based on a Gaussian likelihood and predictive distributions can be deduced from familiar Gaussian formulae. The model incorporates a linear factor structure, which facilitates the use of Gibbs MCMC methods. These features, together with the structure of the priors introduced in Section 4, makes Bayes estimation feasible; indeed, we computed all the results for our benchmark model in a matter of minutes using a 24-core workstation.

Section 5 presents our results for the historical period for which we have data, including rates of convergence and convergence clubs. In Section 6, we turn to the long-run predictive distribution. Section 7 concludes.

2. Data and Descriptive Statistics

2.1 The data

The data are annual values of real per-capita GDP for 113 countries spanning the 118-year period 1900-2017, taken from the Penn World Table (Feenstra, et. al. (2015)) and Bolt et. al.'s (2018) Maddison Project Database. GDP is measured at constant 2011 national prices, expressed in U.S. dollars.¹

¹ Specifically, real GDP is $rgdp^{na}$ from the Penn World Table and population is $population$. We link these series to per-capita GDP $rgdpnipc$ and pop from the Maddison database beginning in the earliest available Penn World Table date for each country (typically 1950).

The 113 countries are those with at least 50 years of available data and 2017 population levels of at least 3 million people. The resulting 113 countries account for 96% of world GDP and 97% of world population in 2017. Of the 69 countries in the Penn World Table that are excluded, 41 are excluded because of limited data (the largest being Ukraine, which has only 38 years of data), 54 because of a small population (the average 2017 population is less than one million for these countries), and 26 for both reasons. The 113 included countries are depicted in Figure 1. The data set is an unbalanced panel with between 36 and 52 countries for the years 1900-1949, 108 countries in 1950, 111 in 1952, and all 113 beginning in 1960. Appendix 1 discusses the data in more detail and lists the countries and sample periods.

The data, in logarithms, are plotted in Figure 2.

2.2 Long-run components

The paths of GDP per capita in Figure 2 exhibit both long-run movements and high frequency fluctuations arising from measurement error, business cycles, and other relatively short-lived sources. Because our interest is in modeling the long-run growth properties of these data, we adopt a procedure that eliminates short-run fluctuations while retaining long-run trends.

In principle, trend extraction can be done using a low-pass filter. The specific method we use is from Müller and Watson (2008, 2018) and reviewed in Müller and Watson (2019). For a given time series y_t , the low-frequency trend \hat{y}_t is the fitted value from the OLS regression of y_t onto a vector X_t which consists of a linear trend, a constant, and $q-1$ low-frequency periodic functions.²

This low-frequency trend extraction method has three useful features. First, as shown in Müller and Watson (2008), it well-approximates an ideal “low-pass” trend that extracts periodicities longer than $2T/q$, where T is the sample size and q is the number of regressors excluding the constant term. We focus on periodicities longer than 14 years, so for countries with a full set of $T=118$ years of data, and we use $q=16 \approx 2 \times 118/14$. Second, as shown in Müller and Watson (2008, 2019), under quite general conditions on the stochastic process for y_t (including unit root and fractionally integrated models), the OLS regression coefficients are approximately

² Müller and Watson (2018) use a constant term and Type II cosine transforms for the periodic regressors to compute the low-frequency trend. Here we also include a linear time trend and, following Müller and Watson (2008), use the $q-1$ eigenvectors of the $T \times T$ covariance matrix of a detrended random walk for the periodic regressors associated with the largest eigenvalues.

jointly asymptotically normal. Thus inference and Bayesian modeling can treat the trend coefficients as Gaussian even if the underlying data are not. Third, this method is in effect a data compression method that reduces the dimensionality of the data from T to $q+1$, which provides computational advantages.

The method is illustrated in Figure 3. Panel (a) focuses on countries with data available over the entire 1900-2017 sample period. The first panel plots the regressors X_t : the linear trend, the constant, and the periodic functions. The remaining panels show the logarithm of GDP per capita for various countries ($y_{i,t}$) and its fitted value $\hat{y}_{i,t}$ from the OLS regression of $y_{i,t}$ onto X_t . The countries are chosen to illustrate how the method extracts the trend component of time series that have quite different historical behavior. In each case, the trend component evidently matches the long run movements in GDP per capita for each country, including fairly subtle shifts such as the multiple pronounced swings in Argentina, the multidecadal slowdown of growth in Germany, the acceleration of growth in India since the 1970s, the shift to slower growth in Mexico following the 1994 Peso crisis, and the plateau in GDP per capita in the United States following the financial crisis recession.

Panels (b) and (c) illustrate the method for countries with data available for only part of the sample. In panel (b) the data are available from 1950-2017, so $T = 68$ and we set $q = 9$ to capture periods longer than 15 years. Even with this shortened sample, the resulting trend component captures the disparate low-frequency patterns in Liberia and Saudi Arabia. Twelve countries have data available over disconnected sub-periods; for example, panel (c) shows data for China, where the data are available from 1929-1938 and then again from 1950-2017. In these cases, the periodic regressors are computed by modifying the method discussed in footnote 2 to accommodate missing values. Details are provided in Appendix 2.

2.3 A first look at the data

Figure 4 plots the low-frequency transformed data for all 113 countries.

We highlight five features of the data that play a role in the econometric model introduced in the next section (for a complementary discussion see Johnson and Papageorgiou (2018, Section II)).

1. Common growth factor. Figures 2 and 4 also show the OECD per-capita level of GDP, computed from the subset of OECD countries available at each date. The OECD aggregate

shows substantial growth over the 118-year sample, increasing 9-fold from \$4.6k in 1900 to \$41.5k in 2017. Average growth for all countries was even greater: the median average annual growth rate for all countries over all available dates was 2.1%, which corresponds to a 12-fold increase on per-capita GDP over 118 years. Despite the evident heterogeneity in growth paths, there is commonality to the growth of the overall cross-country distribution. For example, the average pair-wise correlation of the trends plotted in Figure 4 is 0.58.

2. Variable multidecadal growth rates. As is evident for the eight countries in Figure 3 and as can also be seen by curves for individual countries in Figure 4, growth rates for individual countries have substantial long-run variability. Pritchett (2000) characterized this variability as episodic growth, which led Hausmann, Pritchett, and Rodrik (2005), Jones and Olken (2008), and others to develop empirical models of discrete transitions, or breaks, across growth regimes. As seen in Table 1(a), this variability of long-run growth rates is evident in both developed economies (witness the long-run growth slowdown in the United States over the past two decades) and in non-OECD countries.

3. Cross-section dispersion. Also evident in Figures 2 and 4 is the wide dispersion in the levels of per-capita GDP. This spread is summarized in Table 1(b), which only considers the period for which data on most countries are available (1950-2017). In the cross-country growth literature, convergence in the spread of the log-levels of GDP per capita is referred to as σ -convergence. The cross-sectional standard deviation and the 75%-25% and 90%-10% interquartile ranges show an increase over time, suggesting σ -divergence not σ -convergence. The cross-sectional dispersion has, however, been roughly stable since 1990. In any event, Figures 2 and 4 and Table 1(b) provide no evidence supporting σ -convergence.³

4. Country-specific persistence. Another feature of the data is the extreme persistence of a country's position in the cross-section distribution through time. Quah (1993), Jones (1997, 2016), and Kremer, Onatski and Stock (2001) document this by computing the transition frequencies across different percentiles of the cross-section distribution. Table 2 shows the transition frequencies across quartiles using average per-capita GDP for 1960-88 and 1989-2017. Treating these as Markov transition probabilities, a country in the bottom quartile has more than a 75% chance of remaining in the bottom half of the distribution after 280 years, and similarly

³ We do not engage in formal tests of σ -convergence. Johnson and Papageorgiou (2018) discuss the literature on σ -convergence and of the econometric challenges (power, selection) of tests for σ -convergence.

for a country that starts in the top quartile of the income distribution. For a typical country, the transition across quartiles occurs very slowly.

5. Correlation within groups of countries. The final feature of the data involves the correlation of economic growth within groups of countries. In the econometric model discussed below, groups will be endogenously determined, but these turn out to be related to standard cultural and geographical groupings. Figure 5 gives a visual impression of these correlations using selected 5-country groups with high within-group correlation. We revisit the question of group structure below after describing the econometric model.

These five features of the data – a dominant common factor, variable multidecadal growth rates, relatively constant cross-sectional dispersion, highly persistent relative income levels, and high correlations within groups of countries – are incorporated into the econometric model to which we now turn.

3. A Time Series Model of Cross-country Long-run Growth Dynamics

This section begins by presenting the econometric model, and then discusses its connection to the large empirical growth literature.

3.1 Econometric model

Let $y_{i,t}$ denote the logarithm of per-capita GDP for country i in year t . This section describes a model of the joint stochastic process for $y_{i,t}$ for the 113 countries in our sample. Before providing a detail description of the model we offer a few general remarks about the low-frequency features of the data that the model is designed to capture.

Specifically, two important modeling simplifications follow from our use of the low-frequency transformations of $y_{i,t}$. First, only the low-frequency properties of the stochastic process need to be modeled. In particular, the stationary $I(0)$ dynamics do not need to be modeled because the only feature of those dynamics that enters the joint distribution of the low-frequency components is the $I(0)$ long-run variance. The second simplification follows because the low-frequency properties of the data are summarized by the estimated trend coefficients, which are normally distributed in large samples. Thus, a Gaussian likelihood can be used for low-frequency inference, so that only first two (low-frequency) moments of the process need to be modeled.

While $I(0)$ dynamics are irrelevant over low-frequencies, highly persistent, but stationary, dynamics are relevant. To capture these highly persistent stationary dynamics the model includes components with autocorrelations that decay at the rate ρ^k where ρ is sufficiently close to one that ρ^k is significantly larger than zero even when k is large, say $k = 50, 100$, or even 500 years. Because of their very slow exponential decay, these are called ‘local-to-unity’ AR(1) processes, but it should be understood that the AR(1) label refers only to the low-frequency behavior of the process; general $I(0)$ dynamics are allowed for the shorter-run properties of the process. We will refer to the parameter ρ as the low-frequency AR parameter. For these local-to-unity processes it is also useful to characterize persistence in terms of their *half-life*: for a stationary process x , the half-life is the smallest value of h for which $\text{corr}(x_t, x_{t+h}) = 1/2$, and for a AR(1) process with AR parameter ρ , the half-life solves $\rho^h = 1/2$. Thus, a half-life of $h = 100$ yields $\rho = 0.993$, while $h = 400$ yields $\rho = 0.998$; when ρ is near 1, small changes in ρ lead to large changes in half-life.

The model is designed to capture the five key features of the data evident in the descriptive statistics: (i) long-run global growth, (ii) low-frequency variation in that growth rate, (iii) a roughly stationary distribution of the cross-section around the global growth factor, (iv) highly persistent country-specific deviations from the global factor, and (v) cross-country correlations within groups of countries. We present the model in two steps, focusing first on cross-country covariation and then on temporal covariation.

Cross-country covariation. Common factors are used to capture low-frequency cross-country covariation. The model includes a single common global growth factor, f_t , that affects all countries. The evolution of this factor shifts the entire cross section and is responsible for the time-varying level of per-capita GDP in Figure 2:

$$y_{i,t} = f_t + c_{i,t}, \tag{1}$$

where the country-specific term, $c_{i,t}$, shows country i 's location in the time t cross-section. We use an additional set of factors to describe the cross-country covariance of $c_{i,t}$. There are 35 of these factors, which enter through a clustered hierarchical structure.⁴ Specifically, each country

⁴ Frühwirth-Schnatter and Kaufmann (2008) and Hamilton and Owyang (2012) develop clustered factor structures to explain common business-cycle variability across regions. See Moench, Ng, and Potter (2013) for a discussion and example of hierarchical factor structures.

is allowed to be a member of a single group (or club) whose members share a single common factor. For example, country i might belong to group $J(i)$, with factor $g_{J(i),t}$,

$$c_{i,t} = \mu_c + \lambda_{c,i} g_{J(i),t} + u_{c,i,t} \quad (2)$$

where $u_{c,i,t}$ captures low-frequency variation that is unique to country i , i.e., over low-frequencies, it is uncorrelated across countries. To accommodate low-frequency covariation across countries in different groups, we allow each of the g -level factors to belong to a higher-level group (a group-of-groups), which in turn is affected by a common factor. Thus, the group factor $g_{j,t}$ evolves as

$$g_{j,t} = \lambda_{g,j} h_{K(j),t} + u_{g,j,t} \quad (3)$$

where $K(j)$ denotes j 's group-of-groups, $h_{K(j),t}$ is a common factor for this group, $u_{g,j,t}$ captures idiosyncratic low-frequency variation within group j , and the factor $h_{k,t}$ captures covariation of countries whose groups belong to the k^{th} group-of-groups. The elements of $J(i)$ and $K(j)$ (that is, group and group-of-groups membership) are latent binary indicators so group membership is estimated, not specified *a-priori*. In the empirical model, there are $n_g = 25$ groups and $n_h = 10$ group-of-group factors corresponding to 35 factors, g and h . For symmetry in notation, we denote

$$h_{k,t} = u_{h,k,t}, \quad (4)$$

so that low-frequency variation in $y_{i,t}$ arises from four distinct and uncorrelated sources: f_t , $u_{c,i,t}$, $u_{g,J(i),t}$ and $u_{h,K(J(i)),t}$.

Temporal covariation. The common factor f_t captures common long-run global growth. We model its evolution as an $I(1)$ process (i.e., a low-frequency ‘random walk’) with local drift,

$$\Delta f_t = m_t + \Delta a_t, \quad (5)$$

where Δa_t is $I(0)$ with mean zero and uncorrelated with m_t over low-frequencies. The local growth rate m_t is modeled as a highly persistent AR(1) process with mean μ_m and low-frequency AR coefficient ρ_m . That is,

$$\Delta a_t \sim I(0) \text{ with mean 0 and long-run variance } \sigma_{\Delta a}^2 \quad (6)$$

$$m_t = (1 - \rho_m)\mu_m + \rho_m m_{t-1} + e_t \text{ where } e_t \sim I(0) \text{ with mean 0 and long-run variance } \sigma_{e_m}^2, \quad (7)$$

where the intercept in (7) is written so that the mean of m_t is μ_m .

The model (5)-(7) for the common factor has the following interpretation. If $\sigma_{e_m}^2$ is small relative to $\sigma_{\Delta a}^2$, then f_t evolves over the long-run like a random walk with drift, but with a slowly varying drift term (m_t). If ρ_m were one, then the model would be a low-frequency version of Harvey's (1989) local level model. By specifying ρ_m close to but less than one, the drift term is stochastic but, over a very long horizon, is mean reverting. Thus, the common factor can have persistent excursions in its growth rate – as it evidently has had over the past 118 years (Table 1(a)) – but over the very long run reverts to a mean growth rate μ_m . The persistence of these growth excursions is determined by ρ_m . The variance of m_t over long time spans is $\sigma_m^2 = \sigma_{e_m}^2 / (1 - \rho_m^2)$, which is a key parameter in the model because it determines the magnitude of the persistent growth excursions of the common global factor.

The term $c_{i,t}$ in (1) is the discrepancy between of the log-level of per-capita GDP in country i and the global factor. The descriptive statistics suggest that this is highly persistent. As described above, variation in $c_{i,t}$ arises from the u random variables in (2), (3), and (4). We model each of these variables as stationary but potentially highly persistent. Thus, over the very long run, each country's growth is determined by f_t , but slow mean reversion in $c_{i,t}$ provides country-specific dynamics that are ultimately transitory but may have a half-life of several centuries.

Autoregressive component models are a convenient and flexible way to capture persistence. We therefore model each of the u_t terms as the sum of two independent low-frequency AR(1) processes, say $u_t = u_{1,t} + u_{2,t}$, with low-frequency AR coefficients ρ_1 and ρ_2 and where $(1 - \rho_j L)u_{j,t}$ has long-run variance σ_j^2 for $j = 1, 2$. Different values of $(\rho_1, \rho_2, \sigma_1, \sigma_2)$ allow for processes with, for example, a relatively quickly mean-reverting component (say, a half-life

of 30 years) and very slowly mean-reverting component (say, a half-life of 300 years). In AR models, parameters such as ρ_1 affect both the persistence and variance of the process. To separate persistence from variability we parameterize each u_t process as

$$\begin{aligned} u_t &= \sigma_u \times w_t & (8) \\ w_t &= \zeta w_{1,t} + (1 - \zeta^2)^{1/2} w_{2,t} \\ w_{j,t} &= \rho_j w_{j,t-1} + e_{j,t}, \quad j = 1, 2 \end{aligned}$$

where $w_{1,t}$ and $w_{2,t}$ are independent, each with a unit unconditional variance, and ζ is the weight placed on $w_{1,t}$. In this parameterization, w_t has a unit variance, (ρ_1, ρ_2, ζ) describe the persistence in u_t , and σ_u is its unconditional standard deviation.

There are 148 u -processes, corresponding to the 113 countries in (2), the 25 g -factors in (3) and 10 h -factors in (4). We model each as an independent 2-component low-frequency AR process with its own $(\rho_1, \rho_2, \zeta, \sigma_u)$ parameters.

Relationship of the model with previous work. The model features two forms of β -convergence familiar from the growth literature (c.f., the surveys by Durlauf and Quah (1999) and Johnson and Papageorgiou (2018)). First, in the long run, the expected GDP paths of any two countries i and j are expected to converge in the sense of Bernard and Durlauf (1995, 1996), that is, $\lim_{h \rightarrow \infty} E(y_{i,t+h} - y_{j,t+h} \mid \Omega_t) = 0$, where Ω_t contains the history of y through time t . This convergence obtains in the model because the country-specific terms $c_{i,t+h}$ and $c_{j,t+h}$ exhibit mean reversion to their common mean μ_c , and f_t has the same effect on all countries so that f_t is a single common trend. While all country's forecast paths converge to the same point, the speed of convergence differs across countries because of the heterogeneity in the persistence parameters (ρ_1, ρ_2, ζ) . Said differently, because $c_{i,t}$ is stationary, in this model all countries share a single common trend (f_t), and in this sense are cointegrated. The persistence parameters might be such that this cointegration would not be evident in any century-long sample, however.

Second, in the medium run (which in our model can be a half a century or more), the model also features a form of conditional β -convergence (e.g., Barro (1991), Barro and Sala-i-Martin (1992), and Mankiw, Romer and Weil (1992)) in which $y_{i,t}$ tends toward a growth path with a country-specific level. The vast growth-regression literature has investigated the sources

of heterogeneity in these levels (see Sala-i-Martin (1997) for several examples and Durlauf (2009) for a survey). In our framework, conditional convergence is captured by the AR-component of the various u random variables that determine the evolution of $c_{i,t}$. Each u term is the sum of two independent components, the first with persistence parameter ρ_1 and the second with ρ_2 . If ρ_1 is very close to unity, the first component will be very persistent, for example it can have a half-life of several centuries. While ultimately mean reverting, this component can vary little over, say, 50-year samples and in this sense captures the economic forces underlying the level-shifters included in growth regressions. A smaller value of ρ_2 produces a component with relatively rapid mean reversion; for example, $\rho_2 = 0.98$ produces the 2%-per year convergence rate often found in growth regressions (see Barro (2012)).

The model also features convergence clubs discussed for example in Quah (1996, 1997). Near unit-root dynamics for one of the AR components describing the factors g_t or h_t generates a highly persistent common level component for the group of countries that load on this persistent factor. This persistent group component could have a half-life of several centuries, so that there could be relatively rapid convergence within the club but the club itself converges very slowly to the global factor.

Our model generalizes the model of Raftery et al (2017), which was developed to construct long-horizon forecasts of per-capita global GDP growth as an input to climate change research. That model features a single common factor, proxied by the U.S., which follows a random walk with a drift that breaks in 1973, but is otherwise constant. Their country-specific terms ($c_{i,t}$ in our notation) follow independent zero-mean AR(1) processes. Relative to Raftery et al. (2017), the model here allows for low-frequency variation in the growth rate of the common factor and group convergence dynamics.

There are also notable features that are *not* incorporated in the model. In particular the model does not feature σ -convergence, a narrowing of the cross-sectional distribution over time. In our formulation, while we allow heterogeneity in the variance of $c_{i,t}$ across countries, these variances are constant through time, so the implied variance of the cross-sectional distribution of $c_{i,t}$ is time invariant. This modeling choice is based on the apparent lack of σ -convergence over the 118-year sample shown in Figures 2 and 4. In addition, the model does not incorporate non-stationarities like those postulated in Lucas (2000) and empirically implemented in Startz (2017).

In those models, each country’s growth is governed by a two-state process that determines its convergence to frontier economies: there is no convergence in the first state, but convergence occurs in the second, absorbing state. Following a transition to the convergent state, poor countries grow rapidly and inequality in income levels decreases over time. Long-run point forecasts of future growth from this model may look much like those from the model we implement – both feature unconditional convergence (in expectation) with a rate estimated from historical data – but long-run predictive densities will differ because of the Lucas-Startz framework has σ -convergence whereas ours does not. In addition, compared to those models we allow for (data-influenced) additional variability in the long-run growth rate of the common factor.

4. Bayes Estimation and Prediction

The challenge in specifying a model that describes the joint dynamics of 113 countries is balancing flexibility about the many ways these variables might interact with the limited information in the sample data. The model outlined in the last section strikes one such balance, but at a cost of introducing more than 800 parameters, some of which are only weakly identified by the sample data. With this in mind, we estimate the model using Bayes methods that augment the sample data with judgment about the values of many of these 800+ parameters.

We begin by presenting the priors used in the empirical analysis. These priors are flat (uninformative) about a handful of the model parameters, but are otherwise informative, and therefore require discussion and justification. We then discuss how the computation of the posterior and the predictive distributions takes advantages of the multiple simplifications arising from the use of low-frequency projections combined with the linear factor structure of the model.

4.1 Priors

There are two sets of parameters in the model. One set includes parameters that are common to all countries; this includes the initial condition f_0 , the mean common growth rate μ_m , the persistence parameter ρ_m , the long-run standard deviations $\sigma_{\Delta a}$ and σ_m that characterize the global factor, f_t , in (5) and the parameter μ_c , the common mean of $c_{i,t}$ in (2). The other set of parameters are country- or group-specific; this includes the factor loadings $\{\lambda_{c,i}, \lambda_{g,j}\}$ in (2) and

(3) and the parameters $(\rho_1, \rho_2, \zeta, \sigma_u)$ that describe the evolution of the various u random variables in (2)-(4). We discuss these in turn.

Common parameters. We use uninformative (flat) priors for f_0, μ_m and μ_c , and for $\sigma_{\Delta a}^2$ use a nearly uninformative inverse- χ_1^2 prior that is scaled to have median equal to 0.03^2 .

The prior for (ρ_m, σ_m) is a key informative joint prior governing the long-term distribution of the growth of the common factor. We choose the prior for ρ_m so that the half-life of growth rate excursions (h_m) is roughly a century. Specifically, the prior for ρ_m is such that the half-life $h_m \sim U[50, 150]$, approximated by a grid of 25 discrete values. For σ_m , we specify an independent symmetric triangular informative prior with support $0.1 \leq 100\sigma_m \leq 2.0$, also approximated by a grid of 25 discrete values.

The prior mean for the long-run standard deviation of m_t is 1.05 percentage points of growth. Over the 1900-2017 sample, the mean OECD growth rate was 1.9% so a +/- one (prior) standard deviation range around that mean is 0.9% to 2.9%. This range encompasses the 25-year growth rates for the OECD (and the United States) tabulated in Table 1. The data turn out to be relatively uninformative about the value of σ_m , and long-horizon forecast uncertainty depends on this parameter, so this distribution is a substantive restriction that makes this prior informative for the out-of-sample predictive distributions. We discuss sensitivity of the predictive distributions to this prior in Section 6.

Country- and factor-specific parameters. We use a common framework for these parameters that incorporates an exchangeable prior on a discrete support with a hierarchical structure. Let $\theta_i, i = 1, \dots, m$, denote a set of these parameters, for example, the set of the country-specific factor loadings, $\{\lambda_{c,i}\}$, in (2), so that $m = n$. We specify the common support for θ_i as $\theta_L \leq \theta_i \leq \theta_U$, with values for θ represented by n_θ grid points, $\theta^1, \dots, \theta^{n_\theta}$ between the upper and lower bounds. Given a prior $p = (p_1, \dots, p_{n_\theta})$, the prior distributions for θ_i are i.i.d. with $P(\theta_i = \theta^j) = p_j$, so that the number of θ_i 's taking on the value θ^j has a multinomial distribution. We use a Dirichlet prior with common parameter α/n_θ for the multinomial probabilities, p_j . That is $p \sim D(\alpha/n_\theta)$, where the parameter α is the parameter of a discrete Dirichlet process prior. With the grid points evenly distributed in $[\theta_L, \theta_U]$ and n_θ large, the Dirichlet prior over p with common

parameter α thus shrinks the prior over θ toward an approximately continuous uniform distribution on $[\theta_L, \theta_U]$. Throughout, we use $\alpha = 20$.

This framework has two key features. First, the discrete support for θ_i greatly simplifies the calculations required for the posterior, a point we discuss in more detail below. Second, the hierarchical structure allows the data to inform the posterior for $\{\theta_i\}$ through its effect on the posterior probability assigned to the possible values of θ , that is $P(\theta_i = \theta^j) = p_j$. The Dirichlet prior shrinks these probabilities toward a common value, but as we will see in the empirical analysis, the data modifies this prior in interesting ways. Specifics for each set of parameters are:

- For $\{\lambda_{c,i}\}$ in (2), $\theta_L = 0.0$, $\theta_U = 0.95$, with $n_\theta = 25$ gridpoints evenly distributed between these values. The same prior is used for $\{\lambda_{g,j}\}$ in (3), with an independent Dirichlet process prior.
- The persistence parameters for the various sets of u random variables in (2), (3) and (4) follow independent Dirichlet-multinomial priors. Each u is characterized by (ρ_1, ρ_2, ζ) (see (8)). We found it useful to parameterize the prior for these parameters in terms of the implied half-life for the stochastic process. We specify a joint prior for (ρ_1, ρ_2, ζ) using three independent random variables $U_i \sim U(0,1)$, $i = 1, 2, 3$, and let $\theta = (U_1, U_2, U_3)$. Denote the half-life for w_1 as $h_1 = 25 + 775(U_1)^2$, so that the half-life is between 25 and 800 years, and the implied value of ρ_1 is $\rho_1 = (0.5)^{1/h_1}$. Define ρ_2 similarly using U_2 , and let $\zeta = U_3$. We construct a uniform grid on $\Theta = [0,1]^3$ for (U_1, U_2, U_3) which defines a grid over the values of (ρ_1, ρ_2, ζ) and use $n_\theta = 100$ grid points. A calculation shows that the resulting prior shrinks the half-lives for each u_t toward a distribution with 25th, 50th, and 75th percentiles of 130, 290, and 510 years.
- The prior for the set of scale factors, σ_u in (8) was calibrated relative to a homoskedastic benchmark model. Specifically, let $\{\sigma_{u,c,i}\}$ denote the set of scale factors for the u random variables in (2). We parameterized these as $\sigma_{u,c,i} = s_{c,i} = (1 - \lambda_{c,i}^2)^{1/2} \kappa_{c,i} \omega$, and similarly for $\{\sigma_{u,g,j}\}$, the scale factors for the u -variables in (3), and $\{\sigma_{u,h,k}\}$ in (4). In this parameterization, unit values of $\kappa_{c,i}$, $\kappa_{g,j}$ and $\kappa_{h,k}$ imply that the variance of $c_{i,t}$ is equal to ω^2 for all i . The κ parameters measure the variances of the various components relative to this homoskedastic benchmark. For ω^2 , we use a nearly uninformative inverse- χ_1^2 prior

that is scaled to have median equal to 1. Priors for $\theta = \{\kappa_{c,i}\}, \{\kappa_{g,j}\}$, or $\{\kappa_{h,k}\}$ are independent Dirichlet-multinomial and use $\theta_L = 1/3$, $\theta_U = 3$ and $n_\theta = 25$ evenly spaced grid points.

- The final parameters govern the selection of countries into groups associated with the g -factors in (2) and how these g -factor groups are further grouped using the h -factors in (3). There are $n_g = 25$ g -groups and $n_h = 10$ h -groups (groups-of-groups). Let $t_{c,j,i} = 1$ if country i is a member of group j , and $t_{g,k,j} = 1$ if group j is a member of group-of-group k . We specify these as independent with $P(t_{c,j,i} = 1) = 1/n_g = 1/25$ and $P(t_{g,k,j} = 1) = 1/n_h = 1/10$.

In-sample values of f_t . The final feature of the prior concerns the in-sample values of the common global factor f_t . Since f_t determines the very long-run average growth in our model, it must capture frontier growth, that is growth in the developed economies. To ensure that the in-sample values f_t accord with this interpretation, we force f_t to average growth in developed economies. We do this by imposing a prior on the in-sample values of a population-weighted average of $c_{i,t}$ for OECD countries. Specifically, we assume $\sum_{i \in OECD} c_{i,t} w_i \sim iidN(0, 0.01^2)$, where the w_i weights are proportional to average population of country i over 1965-1974, scaled to sum to one. This prior shrinks the in-sample values of f_t toward the population weighted logarithm of per capita GDP of OECD countries. Since the OECD countries have no missing data after 1950, this means that the value of f_t has very little posterior uncertainty and is effectively treated as observed over the last 65 years. We stress that this constraint is used for the in-sample values of f_t , but not the forecast out-of-sample values.

4.2 Computing the posterior and predictive distributions

Various features of the model provide simplifications for the calculation of posterior. The use of low-frequency projections of the sample data (that is, the OLS regressions of y_t onto X_t from figure 3) yields two: first, using low-frequency projections reduces the effective sample size for each country from T annual observations, to the $q+1$ OLS coefficients. In this application, T is as large as 118 and q is 16.

Second, because these OLS coefficients are low-frequency averages of the sample data with nonrandom weights, the coefficients are approximately jointly normally distributed under quite general conditions. We therefore use a Gaussian likelihood, which allows for analytic posterior calculations for a subset of the model parameters and the use of conditional normal distributions for Gibbs sampling and prediction. Specifically, let Y_i denote the (q_i+1) low-frequency OLS coefficients for country i , where q_i depends on sample size, which differs across countries in our unbalanced sample. Our analysis relies on the data through $Y = (Y_1', Y_2', \dots, Y_n)'$. A central limit result (see Müller and Watson (2019)) yields $Y \stackrel{a}{\sim} N(\mu(\theta), \Sigma(\theta))$ where θ are the mean, long-run variance and persistence parameters of the model in Section 3.1. This normal distribution serves as the likelihood, which together with a prior yields the posterior for the model parameters θ .

Average values of y_t over the forecast period (2018-2027), that is $\bar{y}_{T+1:T+k} = k^{-1} \sum_{i=1}^k y_{T+i}$, are also jointly normally distributed with the sample projection coefficients Y when (k, T) are large, so $\bar{y}_{T+1:T+k} | (Y, \theta) \stackrel{a}{\sim} N(\mu_k(\theta), \Sigma_k(\theta))$. This yields the predictive density for $\bar{y}_{T+1:T+k} | Y$ by averaging the $N(\mu_k(\theta), \Sigma_k(\theta))$ density using the posterior for θ .

A third simplifying feature is the linear factor structure of the model. Conditional on the model parameters, linear Gaussian filtering can be used to generate draws of the factors, which in turn can be used to obtain posterior draws of the parameters. Each of these Gibbs steps is relatively straightforward, only involving low-dimensional multivariate normal random vectors. The corresponding densities are easily computed by simply evaluating the associated quadratic form for any model of time series persistence. (For the original T dimensional data, this would be prohibitively expensive, so instead, one would need to rely on Kalman iterations or other approaches tailored to the assumed form of persistence.)

Fourth, the calculations are simplified by priors that impose a discrete support for many of the parameters. This makes it possible to pre-compute the covariance matrices and their inverses that are the building blocks for the various Gaussian densities used for the likelihood and Gibbs calculations.

Finally, we treat the missing data in our unbalanced panel as missing at random.

Appendix 2 contains a detailed description of the methods we use to sample from the posterior and predictive densities.

5. In-sample Results: Persistence, Cross-country Correlations, and Volatility

The in-sample characteristics of the posterior shed light on various aspects of cross-country growth convergence, including the speed of convergence, the heterogeneity of that convergence, and the existence and composition of convergence clubs. We begin with the global growth factor, then turn to cross-country dynamics.

5.1 Evolution of f_t

Recall that the prior (strongly) shrinks the in-sample values of f_t toward a population-weighted average of $y_{i,t}$ for the OECD countries, so that the resulting in-sample values of f_t roughly coincide with the logarithm of OECD per-capita GDP. This is evident in the first panel in Figure 6, which plots the estimates of f_t along with the trend components of $y_{i,t}$, shaded to differentiate OECD from non-OECD countries. The 67% error bands for f_t computed from the posterior have length of approximately 0.008, with the small value a consequence of the tight in-sample prior.

The evolution of f_t is governed by four parameters, $\sigma_{\Delta a}$, μ_m , σ_m , and ρ_m (see (6) and (7)). Table 3 summarizes the posterior for these parameters. The posterior median for $\sigma_{\Delta a}$ is somewhat lower than a typical estimate for the U.S., but this is consistent with using the OECD average for f_t . The spread of the posterior is roughly what one would find using 16 *i.i.d.* normal observations, that is using $q = 16$ low frequency observations with Δf_t an *i.i.d.* process. The posterior for the average growth rate, μ_m , is centered around 1.9% per year with 67% error bands of roughly $\pm 1\%$ per year. The second panel of Figure 6 shows the posterior estimates for m_t , the local level of Δf_t . The posterior median shows some variation over the sample, but a constant value of 1.9% (its mean) is within the (pointwise) 67% credible set for all dates.

The long-run standard deviation σ_m is an important factor characterizing the evolution of f_t in the out-of-sample forecast period and therefore in determining uncertainty about the future values of y_t . The prior and posterior for σ_m are plotted in Figure 7. The posterior differs little from the prior, so the data have little to say about σ_m at least over the support of the prior. This is also a finding in frequentist inference on the related local-level relative variability parameter (c.f.

Stock and Watson (1997)). The posterior for the persistence in m_t (parameterized using the half-life parameter h_m) is also essentially identical to its prior. (See Table 3).

5.2 Persistence and variability in $c_{i,t}$

The country-specific terms $c_{i,t}$ are functions of $u_{c,i,t}$ in (2) and $u_{g,j,t}$ and $u_{h,k,t}$ for the relevant factors in (3) and (4). Each of these u -terms has its own persistence and variance parameters, so there are many parameters that affect the persistence and variability in each $c_{i,t}$. To summarize these effects, we focus on three characteristics of the marginal distribution of $c_{i,t}$: (i) its long-run standard deviation (σ_c); (ii) its half-life, the value h for which $\text{corr}(c_{i,t}, c_{i,t+h}) = 1/2$; and (iii) the standard deviation of the change in $c_{i,t}$ over a 50-year span, ($\sigma_{(c_{t+50}-c_t)}$). The first two, σ_c and h , are obvious ways to summarize variability and persistence. The third, $\sigma_{(c_{t+50}-c_t)}$, combines both the persistence and long-run variability of $c_{i,t}$ to measure the likely size of long-run (50-year) changes in $c_{i,t}$. For fixed values of σ_c , $\sigma_{(c_{t+50}-c_t)}$ is decreasing in the persistence of the process, while for fixed persistence, it is increasing in σ_c .

The posterior for these parameters is summarized in Table 4. The upper panel shows the posterior pooled over the OECD and non-OECD countries. The posteriors of σ_c and h are plotted in Figure 8. For both OECD and non-OECD countries the country-specific terms, $c_{i,t}$, are highly persistent, but persistence is markedly higher for the OECD countries. The median half-life exceeds 300 years for the pooled OECD countries, but is closer to 200 years for the non-OECD countries. The variance is also smaller for the OECD countries, and taken together, the standard deviation of 50-year changes in $c_{i,t}$ are roughly 1/3 smaller for the OECD countries. Rich countries tend to remain rich, a feature that, in part, defines inclusion in the OECD.

The bottom panel of Table 4 shows results for eight of the 113 countries in the sample. (Results for all countries are given in Appendix 3.) The first six countries are taken from the groups of countries shown previously in Figure 5. Countries that exhibited rapid development show relatively less persistence: for example, Singapore (plotted in panel (b) of Figure 5) has a median half-life of roughly 185 years, compared to, say, the Belgium and the U.S. (shown in panels (e) and (f) of Figure 5) with half-lives of 400 years. Former Soviet-bloc countries exhibit relatively low persistence and large volatility. The country with the highest persistence and lowest variance of 50-year changes is Australia (median half-life = 461 years and $\sigma_{(c_{t+50}-c_t)} = 0.34$

log-points), and the lowest persistence and highest variance country is Liberia (median half-life = 70 years and $\sigma_{(c_{t+50}-c_t)} = 1.34$ log-points).

Figure 9 summarizes the joint posterior for selected features of the $c_{i,t}$ process. Panels (a)-(c) present scatter plots of the posterior means of the initial values $c_{i,0}$ and σ_c (panel (a)), the half-life, h_i (panel (b)), and the 50-year change standard deviation, $\sigma_{(c_{t+50}-c_t)}$ (panel (c)). Countries that were poor at the beginning of the sample (low values of $c_{i,0}$) tend to be more variable, less persistent, and therefore exhibit larger changes over 50-year samples. The lower persistence leads to more rapid convergence towards the global factor for these countries, but the larger variance implies greater uncertainty about their location in the stationary cross-section distribution. Panel (d) highlights the negative relationship between persistence and variability. The final panel shows a negative relationship between volatility (measured by σ_c) and growth (measured by the change $c_{i,t}$ over 1960-2017, the balanced panel sample period); see Ramey and Ramey (1995) for a discussion based on other data.

5.3 Covariance clubs

Correlation between countries in the model arises from three sources: (i) f_t , the global factor, affects all countries, (ii) groups of countries load on the same g -group factor in (2), and (iii) countries might load on different g -factors, but these factors might load on the same h -group-of-group factor in (3). We summarize the resulting pairwise correlations by computing the posterior average population correlations between 50-year changes in $y_{i,t}$ and in c_t (the latter excluding covariability arising from f_t), where again the 50-year horizon is motivated by our interest in long-run covariability.

The average pairwise correlation between 50-year changes in log-per-capita-GDP is 0.59, the largest pairwise correlation is between France and the Netherlands (0.97) and the smallest is between Liberia and Bosnia and Herzegovina (0.29). The average pairwise correlation between the country-specific terms $c_{i,t}$ is, of course, much smaller (0.08); the largest of these is between France and the Netherlands (0.90), and this correlation is less than 0.01 for 38% of the countries.

While the sparse factor structure sets the correlation of 50-year changes in $c_{i,t}$ to zero for many of the country pairs, there are large pairwise correlations between some countries. It turns out that these pairwise correlations between the long-run idiosyncratic growth components

defines natural groups, which we refer to as correlation groups. These correlation groups can be represented as non-directed networks, where the network edges are defined by the strength of the pairwise correlation.

Figure 10 plots the network formed by setting edges (unit values in the undirected adjacency matrix) to be 1 if the posterior mean of the correlation between $c_{i,t}$ and $c_{j,t}$ for countries i and j exceeds 0.25 and = 0 otherwise. The resulting network is fairly sparse with only 11% of the pairwise correlations exceeding 0.25. With this threshold, the network reduces to six separate correlation groups involving three or more countries, displayed separately in panels (a)-(f). Panel (e) shows four additional two-country groups.

It turns out that these groups have natural interpretations that relate to their growth history over the 20th century. The first two groups ((a) and (b)) are Asian countries, where the first group grew rapidly during the 1960-1990 and the second group experienced more recent rapid growth. (Somewhat anomalously, Panama is also included in this latter Asian group.) Group (c) is composed of former Soviet-bloc countries, group (d) are countries with Anglo-Saxon origins, group (e) includes 10 European countries, and group (f) contains 45 low- and middle-income countries, primarily in Africa and Latin America. Panel (g) shows four 2-country groupings, where Austria-Germany and Portugal-Spain have obvious geographic and cultural linkages, and Syria-Yemen suffered coincident civil wars. There are 24 countries that are isolates, not connected to any of these groups.⁵

6. Predictive Distributions

Given the posterior, we computed predictive distributions for country-level GDP growth at 50- and 100-year horizons. We also examined its sensitivity to changes in the priors, checked the calibration of the predictive densities using a pseudo-out-of-sample experiment, and compared the predictive distributions from the multivariate model to the model-implied univariate predictive distributions.

⁵ There is little relationship between the pairwise correlations of $c_{i,t}$ and distance between countries. The simple correlation between the posterior mean of the $c_{i,t}$ - pairwise correlations and distance is -0.2; restricting the sample to countries with pairwise correlations greater than 0.25 leaves this unchanged.

6.1 Baseline predictive distributions

Figure 11 shows 67% and 90% predictive intervals for the global factor f_t , and for eight representative countries. Results for all countries are shown in Appendix 3. The median of the predictive distribution calls for the global factor to increase at an average annual rate of 1.9% from the end of the sample in 2017 through 2118. This translates into a more than 6-fold increase in the level of the global factor. The 67% prediction interval for the average growth rate is quite wide, ranging from 0.9% to 2.7%.

The countries shown in Figure 10 illustrate the range of marginal prediction distributions. The stationarity of $c_{i,t}$ implies that each country tends to mean revert to $f_t + \mu_c$, where μ_c is the mean of $c_{i,t}$ in (2); countries with end-sample-values of $y_{i,t}$ below $f_t + \mu_c$ will tend to grow faster than f_t and similarly for $y_{i,t}$ above $f_t + \mu_c$. The posterior for μ_c is summarized in Table 3; its median is -0.6 with a 67% credible that ranges from -0.8 to -0.5. The rate at which countries converge to this global mean is heterogeneous, so some countries are predicted to converge to the global mean over this 100-year horizon while others do not. For example, the United States is predicted to evolve much like the global factor, albeit with a slightly wider predictive density. Singapore (the second richest country at the end of the sample), is predicted to grow more slowly than average (1.4% per year over the next 100 years) as it mean-reverts down toward the growth path of $f_t + \mu_c$. The end-of-sample values of $y_{i,t}$ for China is near $f_t + \mu_c$, so it is predicted to grow at the same rate as f_t ; this entails a slowdown in its growth rate to that of the global factor. Liberia has very low GDP per capita, high trend variability, and low trend persistence, so it is predicted to revert rapidly to the global mean, however there is great uncertainty about that prediction, and the 90% prediction interval 50 years ahead includes the possibility that its GDP per capita fails to return even to its level in the 1960s.

Figure 12 shows prediction intervals for 50- and 100- year average growth rates for each of the 113 countries, where the countries are ordered from poorest to richest based on end-of-sample values of per-capita GDP. The prediction intervals shift down when moving from the poorest to the richest country, reflecting the mean reversion (convergence) in the model which in turn implies that poor countries are predicted to grow faster than rich countries.

A striking and important feature of the intervals in Figures 11 and 12 is their width, which in all cases exceeds 2 percentage points for 50-year average growth for 67% prediction intervals and is typically 5 percentage points for 90% prediction intervals. For the United States,

for example, the 67% prediction interval for average growth over the next 50 years is 0.6% to 2.7%, and over the next 100 years it is 0.7% to 2.6%.

As can be seen in the figure, the 100-year prediction intervals are narrower than the 50-year intervals. For example, the average width of the 67% bands for $h = 100$ is 2.2 percentage points, but is 2.7 percent points for $h = 50$. Increasing the horizon has two countervailing effects on forecast uncertainty: on the one hand, averaging $I(0)$ processes over longer periods reduces variance, but variances increase for highly persistent processes like those describing m_t , the local level of f . Figure 12 indicates that the first effect dominates, at least for 50- and 100-year forecast horizons.

Table 5 summarizes results for the average growth rates for the groups of countries in Figure 10, using end-of-sample populations to weight the country-specific per-capita values. This weighting scheme suggests that global per-capita income will rise by an annual rate of 2.0% of the next 100 years, resulting a more than seven-fold increase in per-capita GDP. The degree of uncertainty is, however, very wide, with a 67% prediction interval of 1.1% to 3.0% per year. The richer countries are predicted to increase more slowly than the poor countries: the 67% prediction interval for 100-year per-capita GDP growth for non-OECD countries is essentially the same as that for OECD countries, but shifted up by 0.4 percentage points. This pattern of faster growth for the poorer countries also can be seen in the correlation-group breakdowns.

6.2 Sensitivity

We investigated the sensitivity of the model to several key assumptions, three of which we discuss here. The first is the prior distribution for σ_m , the long-run standard deviation of the growth rate trend for f_i in (5). The results are summarized in Table 6.

The parameter σ_m governs the extent to which the local trend growth rate of f_i varies over time, with larger values of σ_m admitting larger variation in the growth rate. Because we treat f_i as effectively observed (the OECD average) within sample, changes in the prior for σ_m have very little effect on the in-sample results on convergence and clubs discussed in Section 5. For the forecasts, however, larger values of σ_m have two important effects. First, larger values of σ_m allows the posterior mean of m_t to vary more and because of the slowdown in OECD growth over the final 25 years of the sample, larger values of σ_m mean that the estimated (filtered) 2017 value of the local growth rate is lower, leading to a lower posterior median growth forecast.

Second, larger values of σ_m allow m_t to vary more over the future, leading to a greater dispersion of growth rates.

The second and third rows of Table 6 summarize the sensitivity of the posterior to changes in the prior for σ_m , specifically shifting the prior in (towards smaller values of σ_m) and out (towards larger values) by 50%. Because the data are largely uninformative about σ_m , changing the prior has a large effect on the posterior for σ_m (Table 6(a)). Because f_t is effectively treated as observed in-sample, $c_{i,t}$ is effectively observed in sample, so changing the prior on the parameters of f_t has essentially no effect on the posterior for the parameters of $c_{i,t}$ (Table 6(b)). When the prior favors smaller values of σ_m , the predicted median growth rate increases and the spread around that median tightens, but when the prior favors larger values of σ_m , the median growth rate falls and the spread widens (Table 6(c)).

We also investigated the sensitivity of the results to q , the number of periodic terms used to obtain the estimated country-level trend for log GDP per capita. A larger value of q includes variation of shorter duration, for example we obtained results using $q = 23$ which corresponds to a low-pass filter keeping variation greater than 10 years. As seen in Table 6, using $q = 23$ increases the estimated variability of a_t (the $I(1)$ term in the evolution of f_t) but results in only small changes in the results about persistence, convergence, and clubs discussed in Section 5. Using $q = 23$ has little effect on the predictive distributions. Using the smaller value $q = 9$, which corresponds to a low-pass cutoff of 26 years, yields results that are very similar to the benchmark model of $q = 16$.

As another check, we re-estimated the model over the 1950-2017 sample, when we have a nearly balanced panel. For these calculations we used $q = 9$, focusing on periods longer than 15 years as in the benchmark specification. The results are shown in the final row of Table 6. The shorter sample suggests a somewhat smaller value for $\sigma_{\Delta a}$ and larger value for σ_m (panel (a)), similar country-specific parameters (panel (b)) and future growth (panel (c)).

6.3 A pseudo-out-of-sample forecasting experiment

Typically, pseudo-out-of-sample (POOS) forecasting experiments are of limited use for evaluating long-horizon forecasts because of the limited number of independent long-horizon POOS time-series observations. However, in our context each of the $n = 113$ countries provides

some independent POOS information about the validity of the predictive distribution. We have carried out a POOS experiment that focuses on this cross-sectional information.

Specifically, in the first experiment we estimated the complete model through time $T_1 = 1977$, and computed joint predictive distributions for the average growth rate of f_t and $y_{i,t}$ for each of the 113 countries over the subsequent $h = 20, 30$ and 40 years. The realized values of $y_{i,t}$ are known over these POOS forecast periods; moreover, the realized value of f_t is well approximated by full-sample estimates $f_{t|T}$ (see Figure 6a). Thus, $c_{i,t|T} = y_{i,t} - f_{t|T}$ provides an accurate estimate for the POOS out-of-sample realized value of $c_{i,t}$. We therefore used $f_{t|T}$ and $c_{i,t|T}$ to evaluate the POOS predictive distributions. Specifically, as is standard for evaluating predictive distributions (see Diebold, Gunther and Tsay (1998)), sample values of the predictive distributions probability integral transform (*PITs*) were computed by evaluating the predictive distributions at the realized POOS values of $f_{t|T}$ and $c_{i,t|T}$. Recall that for a correctly specified predictive distribution, the sample values of the *PIT* are distributed as a $U(0,1)$ random variables.

Table 7 summarizes the resulting *PITs* for the experiment, and two other experiments with used $T_1 = 1987$ (with a forecast horizon $h = 20$ and 30 years) and $T_1 = 1997$ (with a forecast horizon of $h = 20$ years). Table 7 summarizes the results. These result in six forecasts for f_t and with *PIT* values shown in the first column of the table. This is a very small sample of dependent observations, but the *PITs* provide no evidence of misspecification in the predictive distributions for f_t .

There are 113 forecasts $c_{i,t}$ for each POOS experiment and forecast horizon, so these forecasts are more informative about their predictive distributions. The table summarizes the $c_{i,t}$ forecasts by showing the fraction of *PITs* in each of the quartiles or the $U(0,1)$ distribution: ideally, 25% of the *PITs* should be less than 0.25, 25% should be between 0.25 and 0.50, and so forth. The table also shows the Kolmogorov-Smirnov (*KS*) statistics testing the null hypothesis that the actual values were draws from the predictive distributions (i.e., that the *PITs* were $U(0,1)$ draws). The p -values for the *KS* statistic were computed treating the posterior predictive distribution as the truth, which incorporates the model-implied cross-sectional dependence in $c_{i,t}$.

The results suggest that the predictive distributions for $T_1 = 1977$ were somewhat too optimistic: roughly half of the realized values of $c_{i,t}$ lie in the lower quartile of the predictive distributions. The predictive distributions for $T_1 = 1987$ and $T_1 = 1997$ seem to be reasonably well-calibrated.

6.4 Comparison of multivariate forecasts to univariate forecasts

A key feature of the simultaneous model of all countries is that the Bayesian methods have the effect of introducing shrinkage in the parameters. Thus, the forecasts for the individual countries reflect shrinkage to common dynamics. It is thus of interest to compare the forecasts emerging from these joint predictive densities to univariate forecasts that do not use the information from other countries.

We used the joint model described and prior described above to compute univariate forecasts, constructed by treating as missing the data for all countries other than the country at hand. The 67% 50-year ahead forecast intervals are summarized in Figure 13, overlaid on the 50-year intervals in Figure 12. Two features are evident. First, the implied univariate intervals are typically (but not always) substantially wider than the multivariate intervals. Second, the projections for the lower-income countries are typically lower in the univariate than multivariate methods, and in many cases include negative values so that the 67% interval includes 50 years of stagnation or collapse. One driver of this difference, especially for low-income countries, is that the univariate forecasts cannot impose the convergence that is allowed for in the multivariate forecasts, albeit with the possibility that the convergence might be so slow that it is not evident for some countries even in a 100-year sample.

As an illustration, the univariate and multivariate forecast intervals are shown in Figure 14 for selected countries. The univariate forecasts extrapolate country-specific in-sample behavior, so for example the Central African Republic is predicted to continue contracting and India and the Republic of Korea are predicted to continue their rapid growth. Indeed, the median univariate forecasts imply that in 100 years, per-capita GDP in Korea will be more than six times larger than the value in the U.S., and the univariate model produces similarly unreasonable forecasts for other rapidly developing countries. In contrast, for several countries the univariate forecasts are similar to the multivariate forecasts; Denmark and Ecuador, plotted in the figure, are two examples.

7. Concluding Comments

We offer two sets of concluding comments. The first focuses on our empirical application and the second on future applications.

In our application, the data turned out to be informative about many aspects of the analysis. For example, it is clear that there is a wide range of rates of convergence, with some countries having convergence half-lives of less than a century and others having half-lives so long that, in a century-long sample, there is essentially no convergence at all. Similarly, the data support the presence of convergence clubs. In fact, the model delivers convergence clubs as disconnected sub-networks, where those sub-networks have intuitive historical and/or regional interpretations.

One aspect on which the 118 years of data on GDP per capita do not speak strongly is the amount of persistent variation in long-term growth rate of the common factor. The long-run standard deviation, σ_m , is weakly identified in the data. In our model, this weak identification does not substantially influence our in-sample conclusions, such as those about convergence clubs, because we treat the factor f_t as essentially observed in-sample (the OECD mean). But for forecasts 50 and 100 years ahead, the prior on σ_m affects both the mean growth rate of the factor (through the estimate of its long-run growth rate today) and the spread of the predictive distribution. We have proposed a particular prior for the value of σ_m that seems reasonable to us, but others might have different priors. We provided examples of how the predictive distributions would change for alternative candidate priors on σ_m . A virtue of the model is that it reduces a seemingly overwhelming question of what is the future distribution of growth for 113 countries over the next century to a question about a scalar parameter, the relative magnitude of the persistent and nonpersistent changes in the growth rate of the global factor.

Finally, the modeling framework outlined here provided a flexible, yet tractable framework for studying the joint dynamics for a large number of related time series ($n = 113$ countries) over a long span ($T = 118$ years) with data irregularities (missing data). It yielded insights about the joint in-sample behavior of the series and provided sensible long-run joint prediction distributions. This framework holds promise for delivering similar insights in other high dimensional empirical applications involving economic time series.

References

- Barro, R. J. (1991), "Economic Growth in a Cross-Section of Countries," *Quarterly Journal of Economics*, 106(2), pp. 407-443.
- Barro, R.J. (2012), "Convergence and Modernization Revisited," NBER Working Paper 18295.
- Barro, R.J. and X. Sala-i-Martin (1992), "Convergence," *Journal of Political Economy*, 100(2), pp. 223-251.
- Bernard, A. and S.N. Durlauf. (1995), "Convergence of International Output Movements," *Journal of Applied Econometrics*, 10, 97-108.
- Bernard, A. and S.N. Durlauf. (1996), "Interpreting Tests of the Convergence Hypothesis," *Journal of Econometrics*, 71, 161-173.
- Bolt, J., R. Inklarr, H. de Jong, and J.L. van Zanden (2018), "Rebasing 'Maddison': New Income Comparisons and the Shape of Long-Run Economic Development," Groningen Growth and Development Centre Research Memorandum 174.
- Diebold, F.X., T.A. Gunther, and A.S. Tay (1998), "Evaluating Density Forecasts with Applications to Financial Risk Management," *International Economic Review*, 39(4), 863-883.
- Durlauf, S.N. (2009), "The Rise and Fall of Cross-Country Regressions," *History of Political Economy*, 41, pp. 315-333.
- Durlauf, S.N. and D. Quah (1999), "The New Empirics of Economic Growth," in *Handbook of Macroeconomics*, J. Taylor and M. Woodford eds. North-Holland.
- Feenstra, R.C., R. Inklarr, and M.P. Timmer (2015), "The Next Generation of the Penn World Tables," *American Economic Review*, 105(10), pp. 3150-3182.
- Frühwirth-Schnatter and Kaufmann (2008), "Model-based clustering of multiple time series," *Journal of Business and Economic Statistics* 26(1), 2008, 78-89.
- Greenstone, M., E. Kopits, and A. Wolverton (2013), "Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation," *Review of Environmental Economics and Policy* 7(1), 23-46.
- Harvey, A.C., *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge UK: Cambridge University Press, 1989.
- Hamilton, J.D. and M.T. Owyang (2012), "The Propagation of Regional Recessions," *Review of Economics and Statistics*, 94(4), pp. 935-947.

- Hausmann, R., L. Pritchett, and D. Rodrik. (2005). "Growth Accelerations," *Journal of Economic Growth*, 10, 303-329.
- Johnson, P. and C. Papageorgiou (2018), "What Remains of Cross-Country Convergence?" *Journal of Economic Literature*, forthcoming.
- Jones, B.F. and B.A. Olken. (2008). "The Anatomy of Start-Stop Growth," *Review of Economics and Statistics*, 90, 582–587.
- Jones, C.I. (1997), "On the Evolution of the World Income Distribution," *Journal of Economic Perspectives*, 11(3), pp. 19-36.
- Jones, C.I. (2016), "The Facts of Economic Growth," in *Handbook of Macroeconomics*, Vol. 2A, J.B. Taylor and H. Uhlig (eds), pp. 3-68.
- Kremer, M., A. Onatski and J.H. Stock (2001), "Searching for Prosperity," *Carnegie Rochester Conference Series on Public Policy*, 55, pp. 275-303.
- Lucas, R.E. (2000), "Some Macroeconomics for the 21st Century," *Journal of Economic Perspectives*, 14(1), pp. 159-168.
- Maddison, A. (2007). *Contours of the World Economy, 1-2030AD: Essays in Macro-Economic History*, Oxford University Press.
- Mankiw, N.G., D. Romer and D.N. Weil (1992), "A Contribution to the Empirics of Economic Growth," *Quarterly Journal of Economics*, 107(2), pp. 407-437.
- Moench, E., S. Ng and S. Potter (2013), "Dynamic Hierarchical Factor Models," *Review of Economics and Statistics*, 95(5), pp. 1811-1817.
- Müller, U.K. and M.W. Watson (2008), "Testing Models of Low-Frequency Variability," *Econometrica*, Vol. 76, No. 5 (September, 2008), 979–1016.
- Müller, U.K. and M.W. Watson (2016), "Measuring Uncertainty about Long-Run Predictions," *Review of Economic Studies*, 83 (4), October 2016, pp. 1711-1740.
- Müller, U.K. and M.W. Watson (2017), "Low-Frequency Econometrics," in *Advances in Economics and Econometrics Vol. II*, B. Honore, A. Pakes, M. Piazzesi, and L. Samuelson (eds). Cambridge University Press.
- Müller, U.K. and M.W. Watson (2018), "Long-run Covariability," *Econometrica*, Vol. 86, No.3, pp. 775-804.
- Müller, U.K. and M.W. Watson (2019), "Low-Frequency Analysis of Economic Time Series," in preparation for *Handbook of Economics Vol. 7*, S. Durlauf, L.P. Hansen, J.J. Heckman and R. Matzkin (eds). North Holland.

- National Academy of Sciences, Committee on Assessing Approaches to Updating the Social Cost of Carbon, Cropper M.L. and Newell, R.G. co-chairs (2017). *Valuing Climate Changes: Updating Estimation of the Social Cost of Carbon Dioxide*. National Academies Press at <http://www.nap.edu/24651>.
- Pritchett, L. (2000). "Understanding Patterns of Economic Growth: Searching for Hills among Plateaus, Mountains, and Plains," *World Bank Economic Review*, 14, 221-225.
- Quah, D. (1993), "Empirical Cross-Section Dynamics in Economic Growth," *European Economic Review*, 37, pp. 426-434.
- Quah, D. (1996), "Twin Peaks: Growth and Convergence in Models of Distribution Dynamics," *Economic Journal*, 106, pp. 1045-1055.
- Quah, D. (1997), "Empirics for Growth and Distribution" Polarization, Stratification, and Convergence Clubs," *Journal of Economic Growth*, 2(1), pp. 27-59.
- Raftery, A.E., A. Zimmer, D.M.W. Frierson, R. Startz, and P. Liu (2017), "Less than 2° C Warming by 2100 Unlikely," *Nature: Climate Change*, 7(9), pp. 637-641.
- Ramey, G. and V.A. Ramey (1995), "Cross-Country Evidence on the Link Between Volatility and Growth," *American Economic Review*, 85(1), pp. 1138-1151.
- Sala-i-Martin, X. (1997), "I Just Ran Two Million Regressions," *American Economic Review Papers and Proceedings*, 82(2), pp. 178-183.
- Startz, R. (2017), "The Next Hundred Years of Growth: Growth and Convergence," manuscript, U.C. Santa Barbara.
- Stock, J.H. and M.W. Watson (1998), "Asymptotically Median Unbiased Estimation of Coefficient Variation in a Time Varying Parameter Model," *Journal of the American Statistical Association*, 93, pp. 349-358.

Figure 1: Countries included in data set
(shown as shaded areas on map)






Figure 2: GDP per-capita for 113 countries




Notes: The thick black curve is OECD GDP per-capita, computed for those countries in the OECD in 2017 that have data available in the indicated year.

Figure 3: Low-frequency regressors (X_t) and the logarithm of GDP per capita ($y_{i,t}$) and estimated trends ($\hat{y}_{i,t}$) for selected countries

(a) Selected series with data over 1900-2017 (full sample)



(b) Selected series with data over 1950-2017



(c) Series with data over 1929-1938 and 1950-2017





Figure 4: GDP per-capita for 113 countries: Trend components



Notes: The thick black curve is OECD GDP per-capita, computed for those countries in the OECD in 2017 that have data available in the indicated year.

**Figure 5: Selected five-country groups
(low-frequency transformed data)**










Figure 6: In-sample estimates of the global common factor f_t and its local growth rate m_t

(a) Trend components of $y_{i,t}$ for OECD countries (blue), and non-OECD countries (grey) and posterior median of f_t (Black)



(b) Posterior median of m_t and 67% posterior error bands




Figure 7. Prior and posterior for σ_m




Figure 8. Posterior distribution of the long-run standard deviation σ_c and half-life h of $c_{i,t}$ for OECD and non-OECD countries






Figure 9: Cross-country relation between initial conditions, half-life, variability and growth




Notes: Scatterplots are of the posterior means of the indicated parameters, where each observation is a country.

Figure 10. Correlation group networks computed from pairwise correlations of $c_{i,t}$ (average pairwise correlation in parentheses)


(a) Asia-1 (0.65)



(b) Asia-2 (0.51)



(c) Former Soviet-Bloc (0.66)



(d) Anglo-Saxon (0.45)






Figure 10 (Continued)


(e) Europe (0.45)



(f) Other low- and middle-income (0.35)




(g) Pairs



Notes: Pairwise correlations are shown on edges, except in panel (f).


**Figure 11: 67% and 90% prediction intervals for per-capita GDP
Global factor and selected countries**



Notes: Dark shaded areas are 67% prediction intervals; 90% intervals include the light shaded areas. Median forecasts are shown as solid lines. The in-sample value of the global factor and its median forecast are shown in each panel.

**Figure 12: 67% and 90% prediction intervals for average growth rates
Countries ordered from poorest to richest at end sample**

(a) Average growth rate over next 50 years



(b) Average growth rate over next 100 years




Figure 13: Univariate (blue) and Multivariate (red) 67% prediction intervals for average growth rates over next 100 years: Countries ordered from poorest to richest at end sample




Figure 14: Median, 67% and 90% prediction intervals for per-capita GDP: Multivariate (red line and shaded regions) and Univariate (blue line, dashes and dots)




Table 1: Selected summary statistics

(a) Mean growth rates of GDP per capita over 30-year periods (annual percentage growth rates)

	1901-1930	1931-1960	1961-1990	1991-2017
United States	1.4	2.2	2.5	1.5
OECD	1.2	2.0	2.9	1.4
non-OECD	1.4	1.8	2.1	3.0
all	1.3	1.9	2.1	1.9

(b) Cross-sectional distribution of $y_{i,t}$ in selected years (logarithm of GDP per capita)

Average value of $y_{i,t}$ over	Cross-section median	Cross-section standard deviation	Cross-section Inter-quantile range:	
			75%-25%	90%-10%
1950 - 1954	7.8	1.0	1.5	2.5
1971 - 1975	8.3	1.1	1.8	2.9
1992 - 1996	8.6	1.3	2.0	3.5
2013 - 2017	9.3	1.2	2.1	3.3

Table 2: Averages of $y_{i,t}$ over 29 year periods: Fraction of countries of moving from growth quartile i (1960-1988) to quartile j (1989-2017)

		Quartile in 1989-2017			
		1	2	3	4
Quartile in 1960-1988	1	0.79	0.21	0	0
	2	0.21	0.68	0.11	0.00
	3	0	0.11	0.71	0.18
	4	0	0	0.18	0.82

Table 3: Posterior percentiles for selected parameters, f_t and c_t processes

Parameter	Percentiles of Posterior				
	0.05	0.17	0.50	0.84	0.95
(a) Parameters for f_t process (in percentage points)					
$\sigma_{\Delta a}$	2.34	2.70	3.33	4.14	4.87
μ_m	-0.08	0.77	1.85	2.79	3.56
σ_m	0.42	0.73	1.13	1.52	1.76
h_m	50	63	96	133	146
(b) Parameters for c process					
μ_c	-0.97	-0.83	-0.63	-0.45	-0.32

Notes: $\Delta f_t = m_t + \Delta a_t$, where Δa_t is $I(0)$ and m_t is a low-frequency AR(1) with mean μ_m and low-frequency AR coefficient ρ_m . The half-life, h_m , solves $\rho_m^{h_m} = 1/2$.

Table 4: Posterior percentiles for selected parameters, $c_{i,t}$ process

	Half-life			σ_c			$\sigma_{(c_{t+50} - c_t)}$		
	0.17	0.50	0.84	0.17	0.50	0.84	0.17	0.50	0.84
(a) Pooled across countries									
All countries	130	233	389	0.85	1.10	1.37	0.45	0.63	0.86
OECD	201	338	486	0.74	0.91	1.15	0.35	0.44	0.58
Non-OECD	121	211	342	0.91	1.15	1.40	0.52	0.68	0.89
(b) Selected countries									
China	150	222	330	1.00	1.20	1.40	0.59	0.68	0.79
Singapore	111	185	291	0.92	1.13	1.34	0.61	0.72	0.86
Madagascar	148	214	317	1.08	1.24	1.42	0.61	0.72	0.85
Belgium	286	394	507	0.69	0.82	0.98	0.32	0.37	0.44
Russia	78	138	239	0.90	1.10	1.34	0.66	0.79	0.94
United States	303	413	529	0.75	0.89	1.06	0.34	0.39	0.45
Australia	340	461	567	0.71	0.84	1.00	0.30	0.34	0.40
Liberia	50	70	108	1.36	1.51	1.62	1.17	1.34	1.50

Table 5: Percentiles of predictive distributions for average growth over next 50 and 100 years
(2017-population weighted average of country growth rates)

	Percentiles: 50 year horizon			Percentiles: 100 year horizon		
	0.17	0.50	0.84	0.17	0.50	0.84
Global Factor (f_t)	0.92	1.86	2.70	0.92	1.87	2.72
<i>Global aggregates</i>						
All countries	1.03	2.05	3.00	1.06	2.04	2.96
OECD	0.74	1.69	2.62	0.79	1.73	2.62
Non-OECD	1.05	2.13	3.11	1.11	2.10	3.04
<i>Groups of countries from Figure 9</i>						
Correlation group						
Asia-1	0.17	1.47	2.67	0.53	1.62	2.64
Asia-2	0.72	2.01	3.14	0.94	2.01	3.04
Former Soviet Block	0.49	1.80	3.11	0.72	1.82	2.86
Anglo-Saxson	0.62	1.67	2.68	0.70	1.73	2.64
Europe	0.66	1.73	2.77	0.72	1.74	2.72
Other low- and middle-income	1.37	2.43	3.45	1.38	2.38	3.28
Other	0.97	1.94	2.86	1.01	1.93	2.83

Table 6. Sensitivity of selected results to the prior for γ and to the number of low-frequency periodic functions used (q).

(a) Posterior for f process (in percentage points)

Prior for σ_m	q	$\sigma_{\Delta a}$			σ_m		
		Percentiles of posterior			Percentiles of posterior		
		0.17	0.50	0.84	0.17	0.50	0.84
Baseline	16	2.70	3.33	4.14	0.73	1.13	1.52
0.5 × Baseline	16	2.95	3.55	4.3	0.33	0.56	0.76
1.5 × Baseline	16	2.52	3.10	3.84	1.1	1.69	2.29
Baseline	9	1.91	2.91	4.19	0.81	1.29	1.68
Baseline	23	3.40	3.95	4.66	0.65	1.05	1.52
Baseline, 1950-2018	9	1.25	1.76	2.57	1.05	1.37	1.68

(b) Posterior for c process (pooled across all countries)

Prior for σ_m	q	Half-life (years)			$\sigma_{(c_{t+50}-c_t)}$		
		Percentiles of posterior			Percentiles of posterior		
		0.17	0.50	0.84	0.17	0.50	0.84
Baseline	16	130	233	389	0.44	0.63	0.86
0.5 × Baseline	16	129	232	387	0.45	0.63	0.86
1.5 × Baseline	16	130	233	391	0.45	0.63	0.86
Baseline	9	107	209	398	0.43	0.66	0.94
Baseline	23	137	245	395	0.44	0.60	0.82
Baseline, 1950-2018	9	120	229	438	0.39	0.66	0.92

(c) 100-year-ahead predictive distributions for average growth rates (PAAR)

Prior for σ_m	q	Global Factor (f_t)			2014-population weighted average of country growth rates		
		Percentiles of posterior			Percentiles of posterior		
		0.17	0.50	0.84	0.17	0.50	0.84
Baseline	16	0.92	1.87	2.72	1.06	2.04	2.96
0.5 × Baseline	16	1.34	1.97	2.64	1.45	2.17	2.87
1.5 × Baseline	16	0.53	1.75	2.88	0.69	1.95	3.11
Baseline	9	0.65	1.67	2.63	0.85	1.92	2.91
Baseline	23	0.97	1.89	2.81	1.09	2.04	3.01
Baseline, 1950-2018	9	0.86	1.84	2.75	1.03	2.05	3.02

Table 7. Summary of probability integral transforms for pseudo-out-of-sample forecasts

Forecast Horizon	PITs						<i>KS-c (p-value)</i>
	<i>f</i>	Fraction of <i>c</i> -PITs between					
		0 – 0.25	0.25-0.50	0.50-0.75	0.75-1.00		
(a) $T_1 = 1977$							
20 years	0.41	0.57	0.14	0.18	0.11	0.36 (0.01)	
30 years	0.38	0.50	0.15	0.19	0.16	0.28 (0.06)	
40 years	0.30	0.46	0.17	0.17	0.20	0.23 (0.14)	
(b) $T_1 = 1987$							
20 years	0.30	0.30	0.31	0.22	0.17	0.15 (0.42)	
30 years	0.20	0.27	0.26	0.28	0.19	0.09 (0.83)	
(c) $T_1 = 1997$							
20 years	0.24	0.12	0.29	0.29	0.29	0.14 (0.47)	