

## AN ECONOMETRIC MODEL OF INTERNATIONAL GROWTH DYNAMICS FOR LONG-HORIZON FORECASTING

Ulrich K. Müller, James H. Stock, and Mark W. Watson\*

*Abstract*—We develop a Bayesian latent factor model of the joint long-run 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 a sparse correlation pattern (“convergence clubs”) between 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.

### I. Introduction

LONG-RUN planning, policy evaluation, and pricing of long-lived assets require long-horizon forecasts. Issues involved in climate change provide leading examples. For example, 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.

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. The resulting joint predictive distribution can be used to gauge uncertainty about future long-run growth in

individual countries or groupings of countries by region or stage of development.<sup>1</sup> The advantage of such a joint modeling approach over country-specific individual forecasts is not only that one obtains a coherent joint prediction, but it also enables cross-country learning about key growth characteristics and incorporates useful cross-country constraints.

The analysis builds on the long-horizon prediction methods developed in Müller and Watson (2016), but extends that univariate analysis to a large 113-country multivariate framework. We posit a long-run parametric model of international growth dynamics that is informed by the vast empirical literature on international growth, development, and convergence (classic references include Barro, 1991, and Mankiw, Romer, & Weil, 1992; see Jones, 2016, and Johnson & Papageorgiou, 2020, for recent reviews). In particular, the model incorporates five features that the previous literature suggests characterize long-term economic growth. First, the model contains a single global factor to which countries converge in expectation, although the rate of this convergence is allowed to be heterogeneous across countries. Second, if these rates of convergence to the global factor are sufficiently slow, a century-long realization can produce apparent convergence to parallel paths (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 evolves 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 multifactor Bayesian dynamic factor model, where the factors are distinguished by their dynamics and their (latent) commonality across groups of countries.

The focus on long-run dynamics and long-horizon forecasts leads to several simplifications in modeling and estimation. First, it allows us to abstract from short-run and business cycle features by filtering the data to eliminate variation

Received for publication June 3, 2020. Revision accepted for publication September 30, 2020. Editor: Xiaoxia Shi.

\*Müller: Princeton University; Stock: Harvard University and NBER; Watson: Princeton University and NBER.

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. U.M. acknowledges financial support from the National Science Foundation, grant SES-1919336. An earlier version of this paper was titled “An Econometric Model of International Growth Dynamics.”

A supplemental appendix is available online at [https://doi.org/10.1162/rest\\_a\\_00997](https://doi.org/10.1162/rest_a_00997).

<sup>1</sup>The original motivation for this work was the development of long-run probabilistic forecasts of global and regional growth for use in estimating the social cost of carbon, which is the monetized net present value of the economic damages resulting from emitting an additional ton of carbon dioxide. See National Academy of Sciences (2017, chap. 3) for discussion.

associated with periods shorter than fifteen years. With this shorter-run variation eliminated, the model needs only to focus on the longer-run dynamics relevant for long-horizon forecasts. Second, the low-frequency filtering is implemented using weighted averages of the raw data; these low-frequency averages are approximately normally distributed even when the raw data are nonnormal or highly persistent. This allows us to specify a Gaussian probability model for estimation and forecasting, despite the nonnormal characteristics of the underlying data.

We begin in section II with a description of the data, which is a panel of GDP per capita for 113 countries from 1900 to 2017, taken from the Penn World Table (Feenstra, Inklarr, & Timmer, 2015) and the updated Maddison Project Database (Bolt et al., 2018). The panel data set is unbalanced, with missing data for some countries in some years. Plots and descriptive statistics highlight five features of the data, echoing previous findings in the growth literature: 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 III 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. This flexibility comes at the expense of introducing hundreds of unknown parameters.

Section IV takes up the problem of estimating these parameters and computing the long-horizon joint predictive distribution for the 113 countries. We focus on 50- and 100-year-ahead predictions. Bayes estimation of a high-dimensional model ( $n = 113$  countries,  $T = 118$  years and over unknown 800 parameters) with missing data, and with a goal of estimating a joint predictive distribution 100 years into the future, presents considerable 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. And 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 formulas. The model incorporates a linear factor structure, which facilitates missing data and the use of Gibbs Markov chain Monte Carlo (MCMC) methods. These features, together with the structure of the priors introduced in section IV, make Bayes estimation feasible; indeed, we computed all the results for our benchmark model in a matter of minutes using a 24-core workstation.

Section V summarizes results for the historical period for which we have data. These results complement and generalize those found in the empirical growth and convergence literature.

Section VI presents our main results, which are long-horizon (50- and 100-year ahead), joint-predictive distributions for the 113 countries. Results are presented for a baseline specification and several alternatives, including a set of 113 country-specific univariate models. The section also summarizes two external validity exercises: a pseudo-out-of-sample forecasting experiment and an application of the model to long-horizon forecasting for average labor productivity (GDP per worker).

Some concluding remarks are offered in section VII.

## II. Data and Descriptive Statistics

### A. The Data

The data are annual values of real per capita GDP for 113 countries spanning the 118-year period 1900 to 2017, taken from the Penn World Table (Feenstra et al., 2015) and Bolt et al.'s Maddison Project Database (2018). GDP is measured at constant 2011 national prices, expressed in U.S. dollars. (Specifically, real GDP is  $rgdp^{na}$  from the Penn World Table, and population is  $population$ . We link these series to per capita GDP  $rgdpn_{apc}$  and  $pop$  from the Maddison database beginning with the earliest available Penn World Table date for each country—typically 1950.)

The 113 countries are those with at least fifty 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 1 million for these countries), and 26 for both reasons. The data set is an unbalanced panel with between 36 and 52 countries for the years 1900 to 1949, 108 countries in 1950, 111 in 1952, and all 113 beginning in 1960.

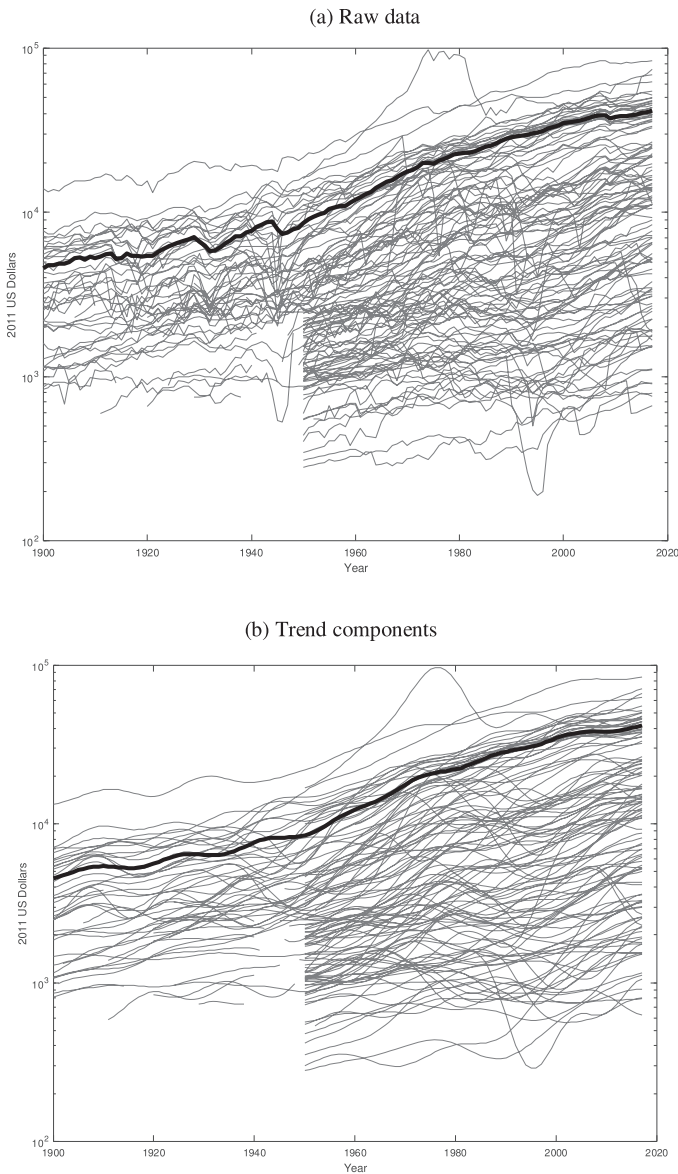
The data, in logarithms, are plotted in figure 1a.

### B. Long-Run Components

The paths of GDP per capita in figure 1a 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 (2020). 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. (Müller & Watson, 2018, use a constant term and type 2 cosine transforms for the periodic regressors

FIGURE 1.—GDP PER CAPITA FOR 113 COUNTRIES



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.

to compute the low-frequency trend. Here we also include a linear time trend and, following Müller & Watson, 2008, use the  $q - 1$  eigenvectors of the covariance matrix of a detrended random walk for the periodic regressors associated with the largest eigenvalues.)

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 filter 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 fourteen years, so for countries with a full set of  $T = 118$  years of data, we use  $q = 16 \approx 2 \times 118/14$ . Second, as shown in Müller and Watson (2008, 2020), under quite general conditions on the stochastic process for  $y_t$  (including unit root and nearly

integrated models), the OLS regression coefficients are approximately jointly 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 considerable computational advantages.

The method is illustrated in figure 2. 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.

Figures 2a and 2b illustrate the method for countries with data available for only part of the sample. In panel b, the data are available from 1950 to 2017, so  $T = 68$  and we set  $q = 9$  to capture periods longer than fifteen 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 subperiods; for example, panel c shows data for China, where the data are available from 1929 to 1938 and then again from 1950 to 2017. In these cases, the periodic regressors are computed by modifying the method discussed above to accommodate missing values. Details are provided in the supplementary material.


### C. A First Look at the Data

Figure 1b plots the low-frequency transformed data for all 113 countries. We highlight five features of the data that are relevant for joint long-horizon forecasts and play a role in the econometric model introduced in the next section.

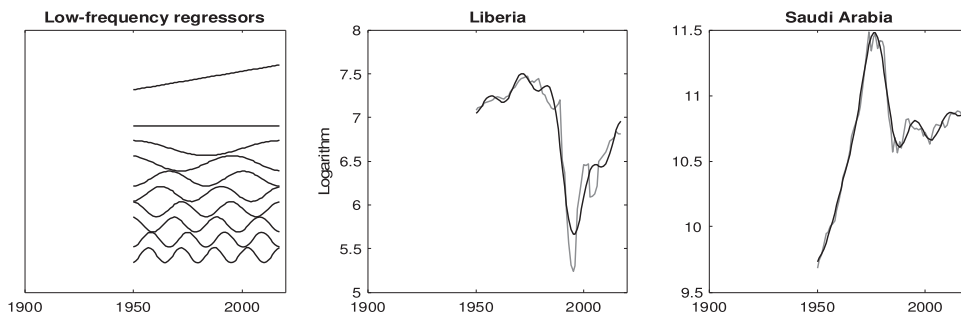
1. *Common growth factor.* Figure 1 shows the OECD per capita level of GDP, computed from the subset of OECD countries available at each date. The OECD aggregate shows substantial growth over the 118-year sample, increasing nine-fold from \$4,600 in 1900 to \$41,500 in 2017. Average growth for all countries was even greater: the median average annual growth rate for all countries over all available dates was 2.1%, which corresponds to a twelve-fold increase of 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

FIGURE 2.—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

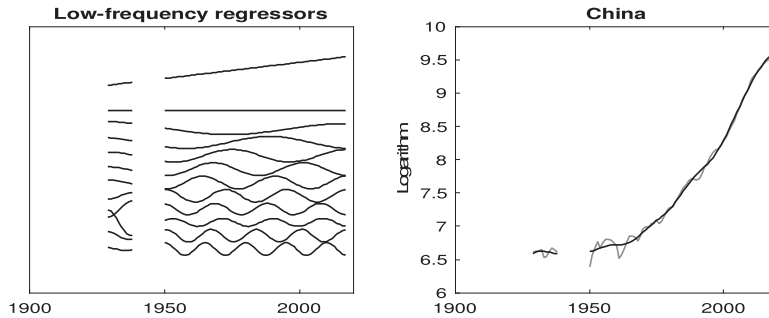
(a) Selected series with data over 1900-2017 (full sample)



(b) Selected series with data over 1950-2017



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



average pair-wise correlation of the trends plotted in figure 1b is 0.58.

2. *Variable multidecadal growth rates.* As is evident for the eight countries in figure 2 and as can also be seen by curves for individual countries in figure 1, growth rates for indi-

vidual 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

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 Interquartile 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

c. 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 1960–1988		Quartile in 1989–2017			
		1	2	3	4
Quartile 1960–1988	1	0.79	0.21	0	0
	2	0.21	0.68	0.11	0
	3	0	0.11	0.71	0.18
	4	0	0	0.18	0.82

in table 1a, 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 non-OECD countries.

3. *Cross-section dispersion.* Also evident in figure 1 is the wide dispersion in the levels of per capita GDP. This spread is summarized in table 1b, which considers only 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  $\sigma$ -convergence. The cross-sectional standard deviation and the 75%-25% and 90%-10% interquartile ranges show an increase over time, suggesting  $\sigma$ -divergence not  $\sigma$ -convergence. The cross-sectional dispersion has, however, been roughly stable since 1990. In any event, figure 1 and table 1b provide no evidence supporting  $\sigma$ -convergence. (Johnson & Papageorgiou, 2018, discuss the literature on  $\sigma$ -convergence and the econometric challenges—power, selection—of tests for  $\sigma$ -convergence.)

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 1c shows the transition frequencies across quartiles using average per capita GDP for 1960 to 1988 and 1989 to 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 the same is true for a country that starts in the top quartile of the in-

come 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 A1 in the supplementary material gives a visual impression of these correlations using selected five-country groups with high within-group correlation.

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 long-horizon joint predictive distributions through the econometric model, to which we now turn.

### III. A Time Series Model of Cross-Country Long-Run Growth Dynamics

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

#### A. 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 detailed 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 the 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  $\rho^k$  where  $\rho$  is sufficiently close to 1 that  $\rho^k$  is significantly larger than 0 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 refer to the parameter  $\rho$  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 an AR(1) process with AR parameter  $\rho$ , 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  $\rho$  is near 1, small changes in  $\rho$  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: long-run global growth, low-frequency variation in that growth rate, a roughly stationary distribution of the cross-section around the global growth factor, highly persistent country-specific deviations from the global factor, and 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 1:

$$y_{i,t} = f_t + c_{i,t}, \quad (1)$$

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  $c_{i,t}$ .

These factors enter through a clustered hierarchical structure.<sup>2</sup> Specifically, each country 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  $i$ ; that is, 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 values of  $J(i)$  and  $K(j)$  (i.e., group and group-of-groups membership) are latent, so group membership is estimated, not specified a priori. 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  $u_{h,K(J(i)),t}$ .

In the empirical model, we allow for a reasonably flexible covariance structure by using  $n_g = 25$  groups and  $n_h = 10$  group-of-group factors corresponding to 35 factors,  $g$  and  $h$ . This hierarchical factor structure, with up to 35 factors and where countries are endogenously and probabilistically assigned to groups, provides a flexible and parsimonious covariance structure for the country-specific components,  $c_{i,t}$ .

*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  $\Delta a_t$  is  $I(0)$  with mean 0 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  $\mu_m$  and low-frequency AR coefficient  $\rho_m$ , that is,

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

$$m_t = (1 - \rho_m)\mu_m + \rho_m m_{t-1} + e_t \text{ where } e_t \sim I(0)$$

$$\text{with mean 0 and long-run variance } \sigma_{e_m}^2, \quad (7)$$

<sup>2</sup>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.

where the intercept in equation (7) is written so that the mean of  $m_t$  is  $\mu_m$ .

The model for the common factor, equations (5) to (7), 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  $\rho_m$  were 1, the model would be a low-frequency version of Harvey's (1989) local-level model. By specifying  $\rho_m$  close to but less than 1, 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 1a), but over the very long run reverts to a mean growth rate  $\mu_m$ . The persistence of these growth excursions is determined by  $\rho_m$ . The variance of  $m_t$  over long time spans is  $\sigma_m^2 = \sigma_{e_m}^2 / (1 - \rho_m^2)$ , 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 equation (1) is the discrepancy between 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 equations (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  $\rho_1$  and  $\rho_2$  and where  $(1 - \rho_j L)u_{j,t}$  has long-run variance  $\sigma_j^2$  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 thirty years) and very slowly mean-reverting component (say, a half-life of 300 years). In AR models, parameters such as  $\rho_j$  affect both the persistence and variance of the process. To separate persistence from variability, we parameterize each  $u_t$  process as

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

where  $w_{1,t}$  and  $w_{2,t}$  are independent, each with a unit unconditional variance, and  $0 \leq \zeta \leq 1$  is the weight placed on  $w_{1,t}$ . In this parameterization,  $w_t$  has a unit variance,  $(\rho_1, \rho_2, \zeta)$  describe the persistence in  $u_t$ , and  $\sigma_u$  is its unconditional standard deviation.

There are 148  $u$ -processes, corresponding to the 113 countries in equation (2), the 25  $g$ -factors in equation (3), and 10  $h$ -factors in equation (4). We model each as an independent two-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  $\beta$ -convergence familiar from the growth literature (cf. the surveys by Durlauf & Quah, 1999, and Johnson & 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,  $\lim_{h \rightarrow \infty} E(y_{i,t+h} - y_{j,t+h} | \Omega_t) = 0$ , where  $\Omega_t$  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  $\mu_c$ , 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  $(\rho_1, \rho_2, \zeta)$ . 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  $\beta$ -convergence (e.g., Barro, 1991; Barro & Sala-i-Martin, 1992; Mankiw et al., 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  $\rho_1$  and the second with  $\rho_2$ . If  $\rho_1$  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, fifty-year samples and in this sense captures the economic forces underlying the level shifters included in growth regressions. A smaller value of  $\rho_2$  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 level component that is common 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 United States, 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  $\sigma$ -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  $\sigma$ -convergence over the 118-year sample shown in figure 1. In addition, the model does not incorporate nonstationarities like those postulated in Lucas (2000) and empirically implemented in Startz (2020). 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 the Lucas-Startz framework has  $\sigma$ -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.

#### IV. 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 section III 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 more than 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 they 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.

##### A. 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  $\mu_m$ , the persistence parameter  $\rho_m$ , the long-run standard deviations  $\sigma_{\Delta a}$  and  $\sigma_m$  that characterize the global factor  $f_t$  in equation (5), and the parameter  $\mu_c$ , the common mean of  $c_{i,t}$  in equation (2). In the other set, the parameters are country or group specific; this includes the factor loadings  $\{\lambda_{c,i}, \lambda_{g,j}\}$  in equations (2) and (3) and the parameters  $(\rho_1, \rho_2, \zeta, \sigma_u)$  that describe the evolution of the various  $u$  random variables in equations (2) to (4). We discuss these in turn.

*Common parameters.* We use uninformative (flat) priors for  $f_0$ ,  $\mu_m$ , and  $\mu_c$ , and for  $\sigma_{\Delta a}^2$  we use a nearly uninformative inverse- $\chi_1^2$  prior that is scaled to have median equal to  $0.03^2$ . We impose a constraint, explained below, that allows  $f_0$  and  $\mu_c$  to be separately identified.

The prior for  $(\rho_m, \sigma_m)$  is a key informative joint prior governing the long-term distribution of the growth of the common factor. We choose the prior for  $\rho_m$  so that the half-life of growth rate excursions ( $h_m$ ) is roughly a century. Specifically, the prior for  $\rho_m$  is such that the half-life  $h_m \sim U[50, 150]$ , approximated by a grid of 25 discrete values. For  $\sigma_m$ , 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  $\pm 1$  (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  $\sigma_m$ , 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 VI.

*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 equation (2), so that  $m = n$ . We specify the common support for  $\theta_i$  as  $\theta_L \leq \theta_i \leq \theta_U$ , with values for  $\theta$  represented by  $n_\theta$  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  $\theta_i$  are i.i.d. with  $P(\theta_i = \theta^j) = p_j$ , so that the number of  $\theta_i$ 's taking on the value  $\theta^j$  has a multinomial distribution. We use a Dirichlet prior with common parameter  $\alpha/n$  for the multinomial probabilities,  $p_j$ . That is,  $p \sim D(\alpha/n_\theta)$ , where the parameter  $\alpha$  is the parameter of a discrete Dirichlet process prior. With the grid points evenly distributed in  $[\theta_L, \theta_U]$  and  $n_\theta$  large, the Dirichlet prior over  $p$  with common parameter  $\alpha$  thus shrinks the prior over  $\theta$  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  $\theta_i$  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  $\theta_i$ , 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 modify this prior in interesting ways. Specifics for each set of parameters are:

- For  $\{\lambda_{c,i}\}$  in equation (2),  $\theta_L = 0.0$ ,  $\theta_U = 0.95$ , with  $n_\theta = 25$  grid points evenly distributed between these values. The same prior is used for  $\{\lambda_{g,j}\}$  in equation (3), with an independent Dirichlet process prior.
- The persistence parameters for the various sets of  $u$  random variables in equations (2), (3), and (4) follow independent Dirichlet-multinomial priors. Each  $u$  is characterized by  $(\rho_1, \rho_2, \zeta)$  (see equation [8]). We specify a joint prior for  $(\rho_1, \rho_2, \zeta)$  in terms of  $\theta = (U_1, U_2, U_3) \in [0, 1]^3$ . Specifically, let the half-life for  $w_1$  be  $h_1 = 25 + 775(U_1)^2$ , so the half-life is between 25 and 80 years and the implied value of  $\rho_1$  is  $\rho_1 = (0.5)^{1/h_1}$ . Define  $\rho_2$  similarly using  $U_2$ , and let  $\zeta = U_3$ . We construct a uniform grid on  $[0, 1]^3$  for  $(U_1, U_2, U_3)$  which defines a grid over the values of  $(\rho_1, \rho_2, \zeta)$  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,  $\sigma_u$  in equation (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 equation (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 equation (3), and  $\{\sigma_{u,h,k}\}$  in equation (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  $\omega^2$  for all  $i$ . The  $\kappa$  parameters measure the variances of the various components relative to this homoskedastic benchmark. For  $\omega^2$ , we use a nearly uninformative inverse- $\chi_1^2$  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 equation (2) and how these  $g$ -factor groups are further grouped using the  $h$ -factors in equation (3). There are  $n_g = 25$   $g$ -groups and  $n_h = 10$   $h$ -groups (groups-of-groups). Let  $\iota_{c,j,i} = 1$  if country  $i$  is a member of group  $j$  and  $\iota_{g,k,j} = 1$  if group  $j$  is a member of group-of-group  $k$ . We specify these as independent with  $P(\iota_{c,j,i} = 1) = 1/n_g = 1/25$  and  $P(\iota_{g,k,j} = 1) = 1/n_h = 1/10$ .

*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  $\sum_{i \in OECD} (y_{i,t} - f_t) w_i^{pop} \sim iidN(0, 0.01^2)$ , where the  $w_i^{pop}$  weights are proportional to average population of country  $i$  over 1965 to 1974, scaled to sum to 1. 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 previous 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. Note that this constraint allows  $f_0$  and  $\mu_c$  to be separately identified.

*B. Computing the Posterior and Predictive Distributions*

Various features of the model provide simplifications for the calculation of the posterior. The use of low-frequency projections of the sample data (i.e., the OLS regressions of  $y_t$  onto  $X_t$  from figure 2) yields two: first, using low-frequency projections reduces the effective sample size for each country from  $T$  annual observations, to the  $q_i + 1 \leq q + 1$  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 & Watson, 2020) yields  $Y \overset{a}{\sim} N(\mu(\gamma), \Sigma(\gamma))$  where  $\gamma$  are the mean, long-run variance and persistence parameters of the model described in section IIIA. This normal distribution serves as the likelihood, which together with a prior yields the posterior for the model parameters  $\gamma$ . Average values of  $y_t$  over the forecast period (2018–2117), 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  $\bar{y}_{T+1:T+k} | (Y, \gamma) \overset{a}{\sim} N(\mu_k(\gamma), \Sigma_k(\gamma))$ . This yields the predictive density for  $\bar{y}_{T+1:T+k} | Y$  by averaging the  $N(\mu_k(\gamma), \Sigma_k(\gamma))$  density using the posterior for  $\gamma$ .

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, involving only 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 slow, 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 precompute 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.

The supplementary material contains a detailed description of the methods we use to sample from the posterior and predictive densities.

## V. In-Sample Results

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 covariance groups. These characteristics affect the long-horizon, joint-predictive distributions that are discussed in the next section. Here we first summarize the in-sample characteristics of the global growth factor and then turn to cross-country dynamics.

### A. 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 3a, 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:  $\sigma_{\Delta a}$ ,  $\mu_m$ ,  $\sigma_m$ , and  $\rho_m$  (see equations [6] and [7]). Table 2a summarizes the posterior for these parameters. The posterior median for  $\sigma_{\Delta a}$  is somewhat lower than a typical estimate for the United States, 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  $\Delta f_t$  an i.i.d. process. The posterior for the average growth rate,  $\mu_m$ , is centered around 1.9% per year with 67% error bands of roughly  $\pm 1\%$  per year. Figure 3b shows the posterior estimates for  $m_t$ , the local level of  $\Delta f_t$ . 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  $\sigma_m$  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  $\sigma_m$  are plotted in figure 3c. The posterior differs little from the prior, so the data have little to say about  $\sigma_m$  at least over the support of the prior. This is also a finding in frequentist inference on the related local-level relative variability parameter (see Stock & 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 2a).

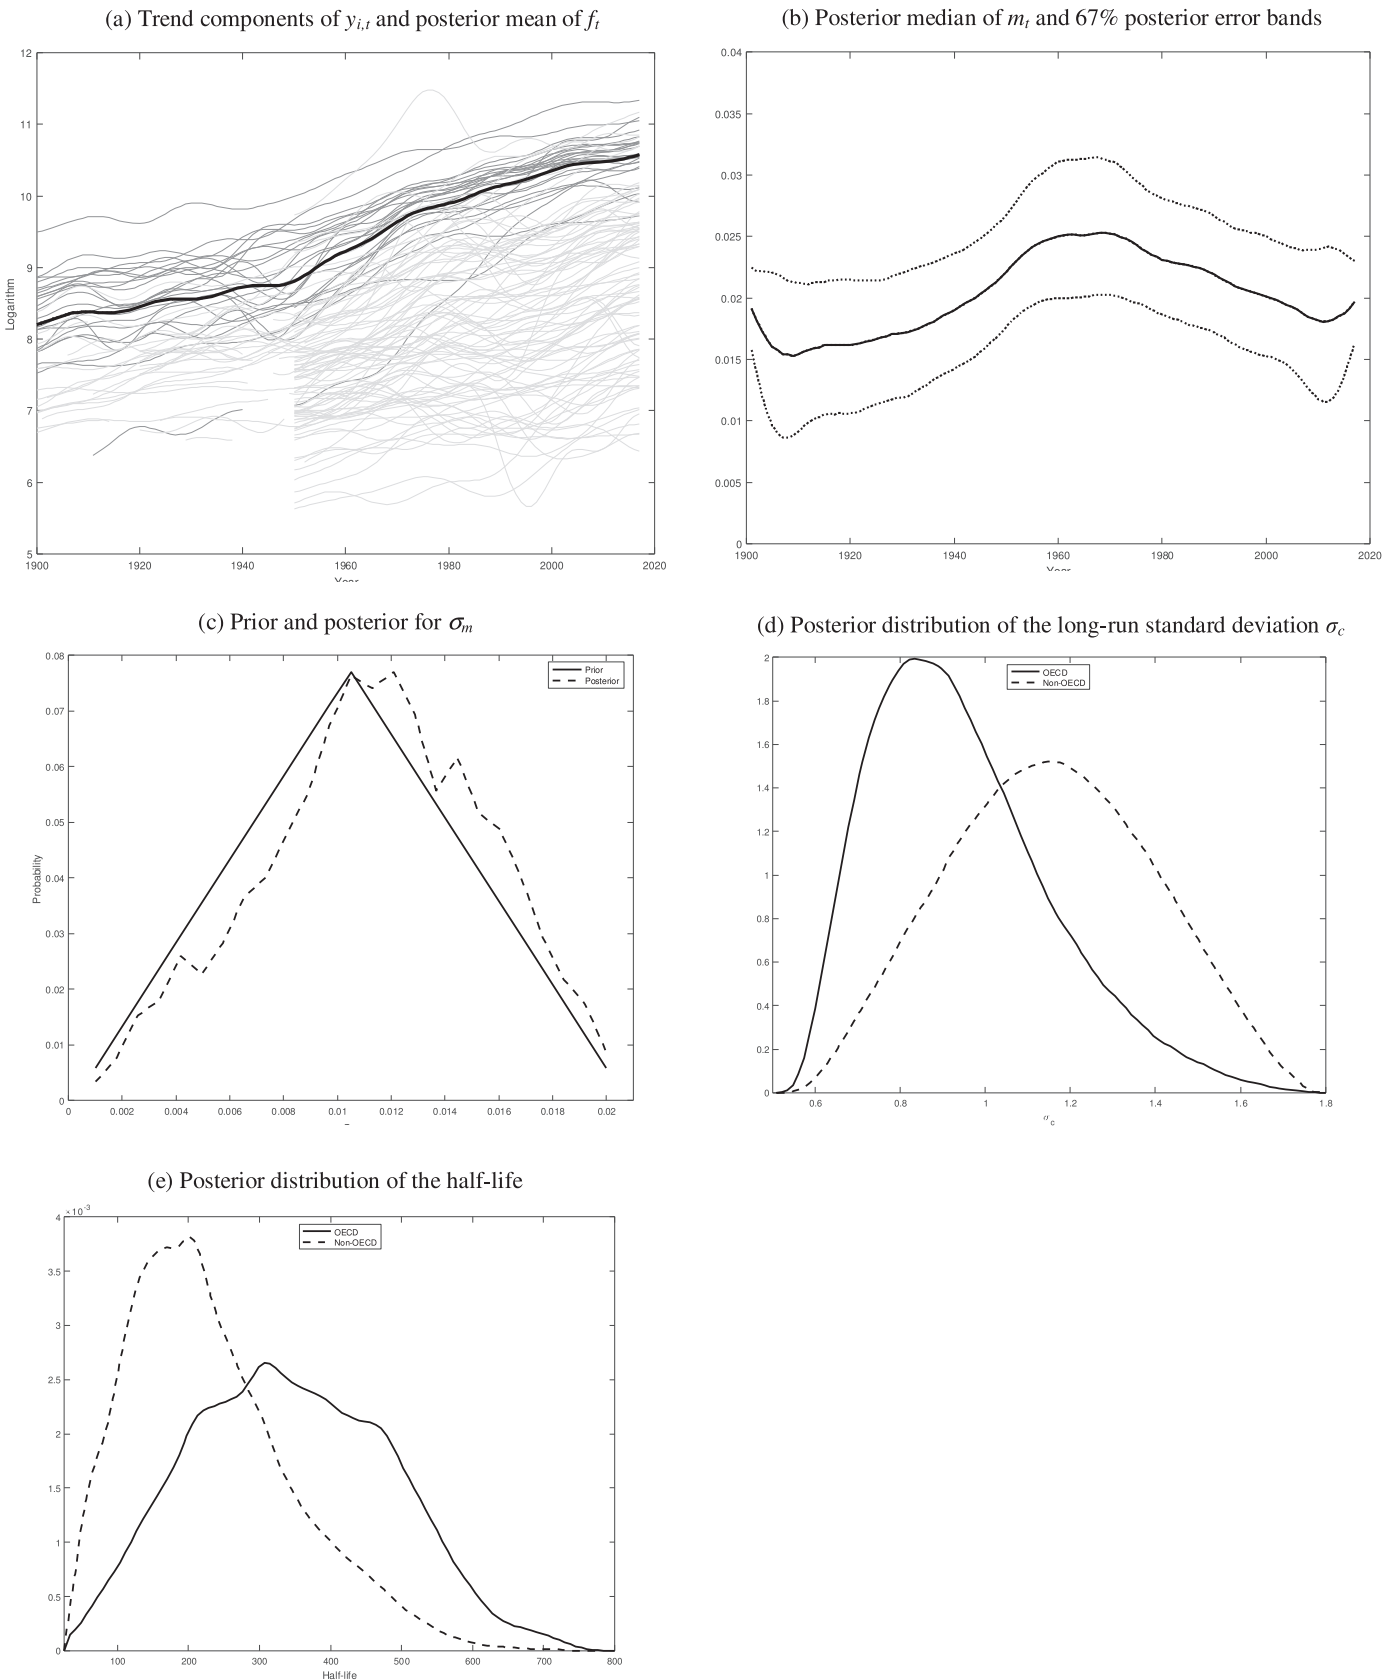
### B. Persistence and Variability in $c_{i,t}$

The country-specific terms  $c_{i,t}$  are functions of  $u_{c,i,t}$  in equation (2) and  $u_{g,j,t}$  and  $u_{h,k,t}$  for the relevant factors in equations (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  $c_{i,t}$ : (a) its long-run standard deviation ( $\sigma_c$ ); (b) its half-life, the value  $h$  for which  $\text{corr}(c_{i,t}, c_{i,t+h}) = 1/2$ ; and (c) the standard deviation of the change in  $c_{i,t}$  over a fifty-year span ( $\sigma_{(c_{t+50}-c_t)}$ ). The first two,  $\sigma_c$  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 (fifty-year) changes in  $c_{i,t}$ . For fixed values of  $\sigma_c$ ,  $\sigma_{(c_{t+50}-c_t)}$  is decreasing in the persistence of the process, while for fixed persistence, it is increasing in  $\sigma_c$ .

The posterior for these parameters is summarized in table 2b. The upper panel shows the posterior pooled over the OECD and non-OECD countries. The posteriors of  $\sigma_c$  and  $h$  are plotted in figures 3d and 3e. 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 deviations of fifty-year changes in  $c_{i,t}$  are roughly one-third 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 2b shows results for 8 of the 113 countries in the sample. (Results for all countries are given in the supplementary material.) The first six countries are taken from the groups of countries shown in figure A1 in the supplementary material. Countries that exhibited rapid development show relatively less persistence; for example, Singapore has a median half-life of roughly 185 years compared to, say, Belgium and the United States 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

FIGURE 3.—SELECTED FEATURES OF POSTERIOR



In panel a,  $y_{i,t}$  is plotted as darker curves for OECD countries.

TABLE 2.—POSTERIOR PERCENTILES

a. 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
i. Parameters for $f_t$ process (in percentage points)					
$\sigma_{\Delta a}$	2.34	2.70	3.33	4.14	4.87
$\mu_m$	-0.08	0.77	1.85	2.79	3.56
$\sigma_m$	0.42	0.73	1.13	1.52	1.76
$h_m$	50	63	96	133	146
ii. Parameters for $c$ process					
$\mu_c$	-0.97	-0.83	-0.63	-0.45	-0.32

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

b. Posterior percentiles for selected parameters, $c_{i,t}$ process									
	Half-life			Percentiles of Posterior			$\sigma_{(c_{t+50}-c_t)}$		
	0.17	0.50	0.84	0.17	0.50	0.84	0.17	0.50	0.84
i. 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
ii. 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

(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 A2 in the supplementary material summarizes the joint posterior for selected features of the  $c_{i,t}$  process. It shows that countries that were poor at the beginning of the sample (low values of  $c_{i,0}$ ) tend to be more variable and less persistent, and therefore they exhibit larger changes over fifty-year samples. The lower persistence leads to more rapid convergence toward the global factor for these countries, but the larger variance implies greater uncertainty about their location in the stationary cross-section distribution. The posterior also shows a negative relationship between persistence and variability and between volatility and growth over the sample period. Ramey and Ramey (1995) provide discussion based on other data.

### C. Correlation between Countries

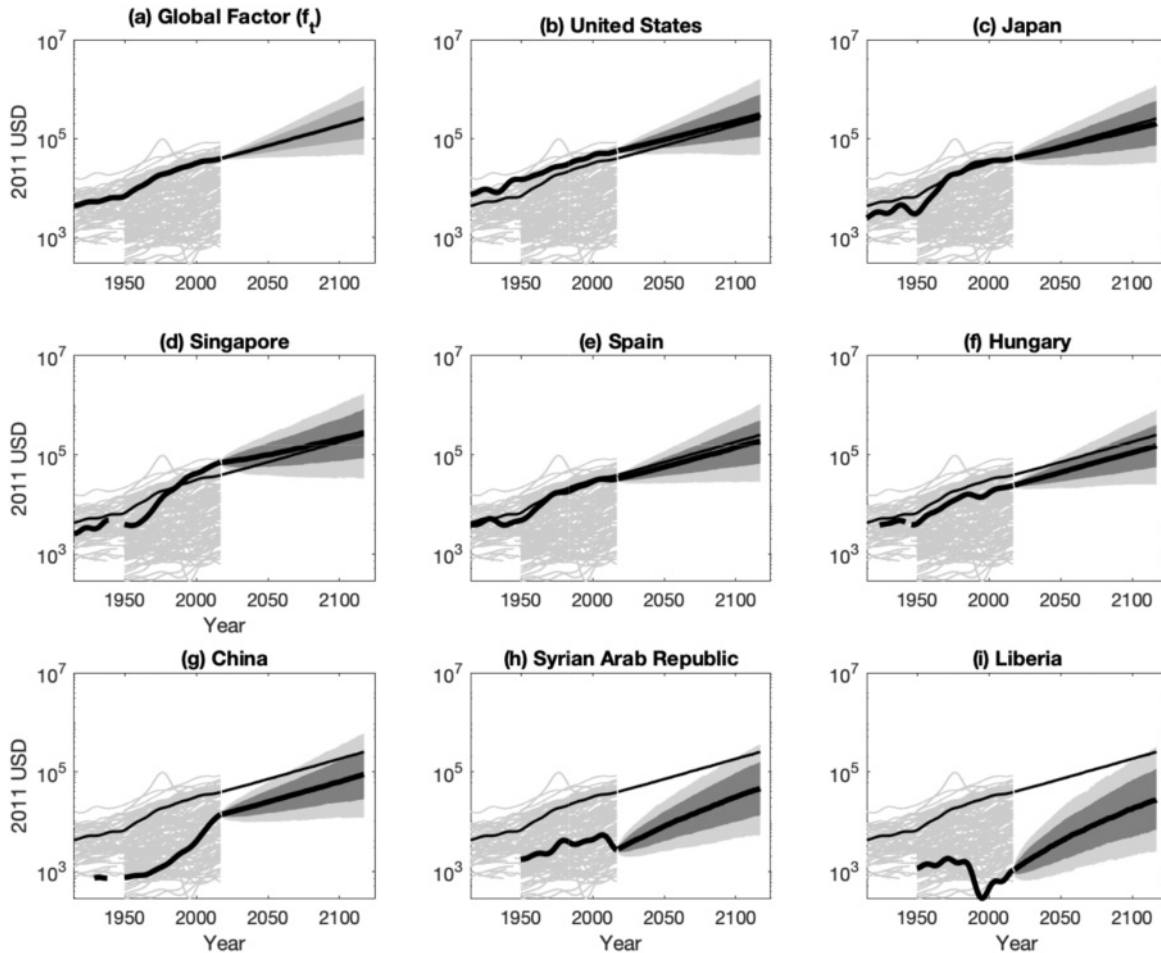
Correlation between countries in the model arises from three sources: (a)  $f_t$ , the global factor, affects all countries; (b) groups of countries load on the same  $g$ -group factor in equation (2); and (c) countries might load on different

$g$ -factors, but these factors might load on the same  $h$ -group-of-group factor in equation (3). We summarize the resulting pairwise correlations by computing the posterior average population correlations between fifty-year changes in  $y_{i,t}$  and in  $c_t$  (the latter excluding covariability arising from  $f_t$ ), where again the fifty-year horizon is motivated by our interest in long-run covariability.

The average pairwise correlation between fifty-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 country pairs.

In many cases, large pairwise correlations are associated with familiar groupings of countries. For example, one grouping includes the early rapid-developing Asian countries (Hong Kong, Korea, Malaysia, Singapore, Taiwan, and Thailand), with an average pairwise of 0.65 for  $c_{i,t}$ . Another includes the former Soviet-bloc countries of Bulgaria, Croatia, Hungary, Romania, Russia, and Serbia, with an average pairwise correlation of 0.66; and yet another includes the

FIGURE 4.—67% AND 90% PREDICTION INTERVALS FOR PER CAPITA GDP GLOBAL FACTOR AND SELECTED COUNTRIES



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.

Anglo-Saxon countries Australia, Canada, New Zealand, the United States, and the United Kingdom, with an average pairwise correlation of 0.45.

Pairwise correlations for all countries are given in the supplementary material.

## VI. Predictive Distributions

This section summarizes the main findings of the paper: the long-horizon predictive distribution of GDP per capita for the 113 countries in the sample and various groupings of these countries. Predictive distributions are shown for 50- and 100-year horizons. The section also discusses sensitivity of the forecasts to changes in the priors, summarizes a pseudo-out-of-sample experiment that checks the calibration of predictive distributions, compares the predictive distributions from the multivariate model to the model-implied univariate predictive distributions, and repeats the analysis using the same model and priors to compute predictive distributions for average labor productivity (GDP per worker) instead of GDP per capita.

### A. Baseline Predictive Distributions

Figure 4 shows 67% and 90% predictive intervals for the global factor  $f_t$  and for eight representative countries. Results for all countries are shown in the supplementary material. 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 six-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 4 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  $\mu_c$  is the mean of  $c_{i,t}$  in equation (2); countries with end-sample values of  $y_{i,t}$  below  $f_t + \mu_c$  tend to grow faster than  $f_t$  and similarly for  $y_{i,t}$  above  $f_t + \mu_c$ . The posterior for  $\mu_c$  is summarized in table 2a.ii; its median is  $-0.6$  with a 67% credibility 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 are 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 fifty years ahead includes the possibility that its GDP per capita fails to return even to its level in the 1960s.

Figure 5 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 ones.

A striking and important feature of the intervals in figure 5 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%.

While prediction intervals for the level of per capita GDP increase with forecast horizon (see figure 4), the 100-year prediction intervals for average growth rates are narrower than the 50-year intervals (see figure 5). For example, the average width of the 67% bands for  $h = 100$  is 2.2 percentage points, but it is 2.7 percentage points for  $h = 50$ . Increasing the horizon has two countervailing effects on forecast uncertainty for average growth rates: averaging  $I(0)$  processes over longer periods reduces variance, while variances increase for averages of highly persistent processes like those describing  $m_t$ , the local level of  $f$ . Figure 5 indicates that the first effect dominates, at least for 50- and 100-year forecast horizons.

Table 3 summarizes results for various groupings of countries 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% during the next 100 years, resulting in 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 grow 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 country groupings used in the International Monetary Fund's

World Economic Outlook (2020), where the median of the 100-year-ahead predictive distributions calls for an average annual growth rate of 2.6% for sub-Saharan Africa and 1.7% for the advanced economies.

### B. 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  $\sigma_m$ , the long-run standard deviation of the growth rate trend for  $f_t$  in equation (5). The results are summarized in table 4.

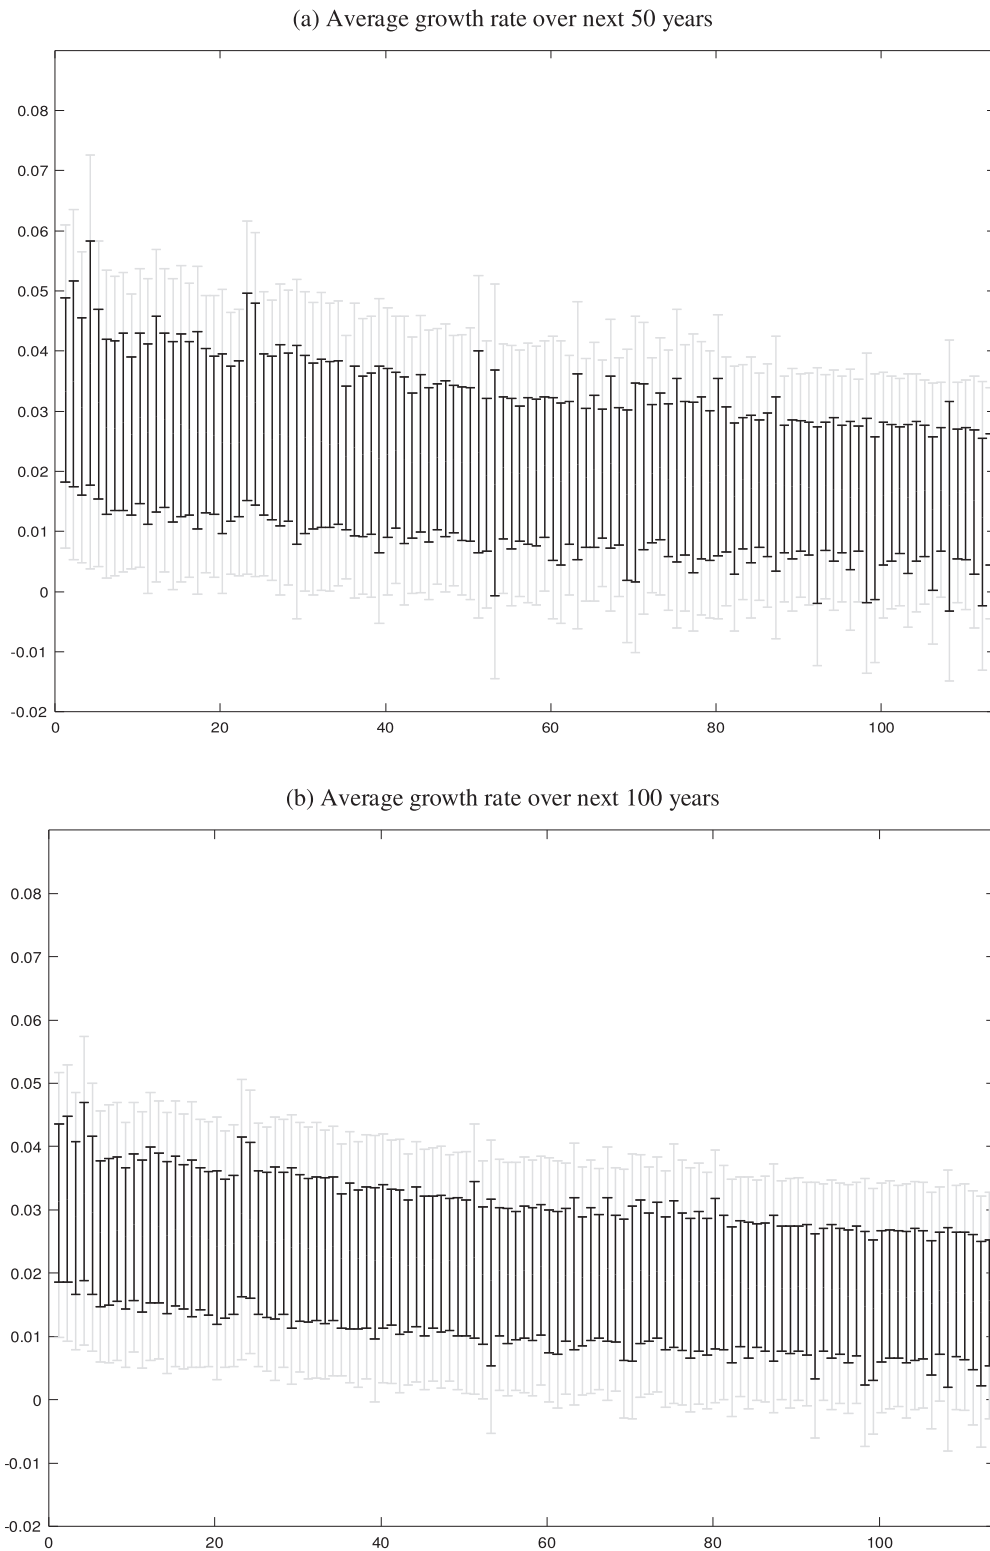
The parameter  $\sigma_m$  governs the extent to which the local trend growth rate of  $f_t$  varies over time, with larger values of  $\sigma_m$  admitting larger variation in the growth rate. Because we treat  $f_t$  as effectively observed (the OECD average) within sample, changes in the prior for  $\sigma_m$  have very little effect on the in-sample results on convergence and clubs discussed in section V. For the forecasts, however, larger values of  $\sigma_m$  have two important effects. First, larger values of  $\sigma_m$  allow 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  $\sigma_m$  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  $\sigma_m$  allow  $m_t$  to vary more over the future, leading to a greater dispersion of growth rates.

The second and third rows of table 4 summarize the sensitivity of the posterior to changes in the prior for  $\sigma_m$ , specifically shifting the prior in (toward smaller values of  $\sigma_m$ ) and out (toward larger values) by 50%. Because the data are largely uninformative about  $\sigma_m$ , changing the prior has a large effect on the posterior for  $\sigma_m$  (table 4a). Because  $f_t$  is effectively treated as observed in-sample, so is  $c_{i,t}$ , 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 4b). When the prior favors smaller values of  $\sigma_m$ , the predicted median growth rate increases and the spread around that median tightens, but when the prior favors larger values of  $\sigma_m$ , the median growth rate falls and the spread widens (table 4c).

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 that extracts periodicities longer than ten years. As seen in table 4, 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 V. Using  $q = 23$  has little effect on the predictive distributions. Using the smaller value of  $q = 9$ , 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 reestimated the model over the 1950–2017 sample, when we have a nearly balanced panel. For

FIGURE 5.—67% (BLACK) AND 90% (GRAY) PREDICTION INTERVALS FOR AVERAGE GROWTH RATES: COUNTRIES ORDERED FROM POOREST TO RICHEST AT END SAMPLE



these calculations, we used  $q = 9$ , focusing on periods longer than fifteen years as in the benchmark specification. As can be seen in table 4, the shorter sample suggests a somewhat

smaller value for  $\sigma_{\Delta a}$  and larger value for  $\sigma_m$  (panel a), similar country-specific parameters (panel b), and future growth (panel c).

TABLE 3.—PERCENTILES OF PREDICTIVE DISTRIBUTIONS FOR AVERAGE GROWTH OVER NEXT 50 AND 100 YEARS: POPULATION WEIGHTED AVERAGE OF COUNTRY GROWTH RATES, 2017

	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
Global aggregates						
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
Selected IMF-WEO groupings						
Advanced economies	0.69	1.64	2.57	0.74	1.68	2.58
Euro area	0.65	1.68	2.66	0.74	1.72	2.64
G7	0.69	1.65	2.61	0.75	1.69	2.60
Emerging and developing economies	1.06	2.13	3.11	1.11	2.11	3.04
Emerging and developing Asia	0.80	2.02	3.12	0.97	2.03	3.02
ASEAN-5	0.90	1.98	3.03	1.00	2.01	2.97
Emerging and developing Europe	0.65	1.79	2.87	0.78	1.80	2.75
Latin America and Caribbean	0.99	2.04	2.99	1.03	2.03	2.93
Middle East and Central Asia	1.15	2.17	3.13	1.19	2.15	3.06
Sub-Saharan Africa	1.51	2.63	3.71	1.51	2.55	3.50

The country groups shown in the bottom panel are from the IMF's World Economic Outlook (2020).

### C. Forecasts for Average Labor Productivity

Thus far, the focus has been on forecasting per capita values of GDP ( $Y/Pop$ ). A related exercise focuses instead on average labor productivity ( $Y/L$ ). Employment data are available in the Penn World Table (PWT) but not the Maddison Project Database, so the sample period is restricted to 1950 to 2017. We used these data and the model of section III to estimate the posterior and long-horizon predictive distribution for average labor productivity.

The supplementary material contains detailed results. The final row in each panel of table 4 summarizes a few key results. The posteriors for the parameters using  $Y/L$  are similar to those using  $Y/Pop$  (panels a and b of table 4), while forecasts are for slightly slower growth and more uncertainty (panel c of table 4).

### D. 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 3a). 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, & 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 A4 in the supplementary material summarizes the resulting PITs for the experiment, and two other experiments use  $T_1 = 1987$  (with a forecast horizon  $h = 20$  and  $30$  years) and  $T_1 = 1997$  (with a forecast horizon of  $h = 20$  years). 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 PITs from these forecasts are summarized in table A4. 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.

### E. 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



TABLE 4.—SENSITIVITY OF SELECTED RESULTS TO THE PRIOR FOR  $\sigma_m$ , THE NUMBER OF LOW-FREQUENCY PERIODIC FUNCTIONS USED ( $q$ ), AND SAMPLE PERIOD

a. Posterior for $f$ process (in percentage points)								
Prior for $\sigma_m$	$q$	Start Date	$\sigma_{\Delta a}$			$\sigma_m$		
			Percentiles of posterior			Percentiles of posterior		
			0.17	0.50	0.84	0.17	0.50	0.84
Baseline	16	1900	2.70	3.33	4.14	0.73	1.13	1.52
0.5 × Baseline	16	1900	2.95	3.55	4.3	0.33	0.56	0.76
1.5 × Baseline	16	1900	2.52	3.10	3.84	1.1	1.69	2.29
Baseline	9	1900	1.91	2.91	4.19	0.81	1.29	1.68
Baseline	23	1900	3.40	3.95	4.66	0.65	1.05	1.52
Baseline	9	1950	1.25	1.76	2.57	1.05	1.37	1.68
Using $Y/L$ instead of $Y/Pop$								
Baseline, 1950–2017	9	1950	1.22	1.78	2.71	1.13	1.45	1.76

b. Posterior for $c$ process (pooled across all countries)								
Prior for $\sigma_m$	$q$	Start Date	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	1900	130	233	389	0.44	0.63	0.86
0.5 × Baseline	16	1900	129	232	387	0.45	0.63	0.86
1.5 × Baseline	16	1900	130	233	391	0.45	0.63	0.86
Baseline	9	1900	107	209	398	0.43	0.66	0.94
Baseline	23	1900	137	245	395	0.44	0.60	0.82
Baseline	9	1950	120	229	438	0.39	0.66	0.92
Using $Y/L$ instead of $Y/Pop$								
Baseline, 1950–2017	9	1950	119	252	479	0.35	0.61	0.91

c. 100-year-ahead predictive distributions for average growth rates (PAAR)								
Prior for $\sigma_m$	$q$	Start date	Global factor ( $f_t$ )			2017-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	1900	0.92	1.87	2.72	1.06	2.04	2.96
0.5 × Baseline	16	1900	1.34	1.97	2.64	1.45	2.17	2.87
1.5 × Baseline	16	1900	0.53	1.75	2.88	0.69	1.95	3.11
Baseline	9	1900	0.65	1.67	2.63	0.85	1.92	2.91
Baseline	23	1900	0.97	1.89	2.81	1.09	2.04	3.01
Baseline	9	1950	0.86	1.84	2.75	1.03	2.05	3.02
Using $Y/L$ instead of $Y/Pop$								
Baseline	9	1950	0.65	1.72	2.70	0.78	1.88	2.96


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% fifty-year-ahead univariate forecast intervals are overlaid on multivariate intervals in figure 6a. 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 fifty years of stagnation or collapse. These differences obtain in spite of the common model and priors used in both the multivariate and univariate models. One driver of this difference 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 6b for selected countries. The

FIGURE 6.—UNIVARIATE AND MULTIVARIATE PREDICTIONS



(a) Univariate (gray) and multivariate (black) 67% prediction intervals for average growth rates over the next 100 years. Countries are ordered from poorest to richest at end sample. (b) Median, 67% and 90% prediction intervals for per capita GDP: Multivariate (line and shaded regions) and univariate (dashes and dots).

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 United States, 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.

## VII. Concluding Remarks

We offer three sets of concluding comments. The first two focus on our empirical application and the third on future applications.

First, our model contains many parameters relative to the information in the sample, and this raises a concern about overfitting. But the use of informative priors, such as those used in our application, helps guard against overfitting. And the (admittedly limited) pseudo-out-sample forecasting experiment and application using the same model for labor productivity provides some comfort about overfitting.

Second, 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 are consistent with a sparse long-run correlation pattern, that is, “convergence clubs.”

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,  $\sigma_m$ , 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  $\sigma_m$  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  $\sigma_m$  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  $\sigma_m$ . A virtue of the model is that it reduces a seemingly overwhelming question of what the future distribution of growth is 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 provides a flexible, yet tractable structure 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 yields 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., “Economic Growth in a Cross-Section of Countries,” *Quarterly Journal of Economics* 106:2 (1991), 407–443. 10.2307/2937943

- “Convergence and Modernization Revisited,” NBER working paper 18295 (2012).
- Barro, R. J., and X. Sala-i-Martin, “Convergence,” *Journal of Political Economy* 100:2 (1992), 223–251. 10.1086/261816
- Bernard, A., and S. N. Durlau, “Convergence of International Output Movements,” *Journal of Applied Econometrics* 10 (1995), 97–108. 10.1002/jae.3950100202
- “Interpreting Tests of the Convergence Hypothesis,” *Journal of Econometrics* 71 (1996), 161–173. 10.1016/0304-4076(94)01699-2
- Bolt, J., R. Inklarr, H. de Jong, and J. L. van Zanden, “Rebasing ‘Maddison’: New Income Comparisons and the Shape of Long-Run Economic Development,” Groningen Growth and Development Centre research memorandum 174 (2018).
- Diebold, F. X., T. A. Gunther, and A. S. Tay, “Evaluating Density Forecasts with Applications to Financial Risk Management,” *International Economic Review* 39:4 (1998), 863–883. 10.2307/2527342
- Durlauf, S. N., “The Rise and Fall of Cross-Country Regressions,” *History of Political Economy* 41 (2009), 315–333. 10.1215/00182702-2009-030
- Durlauf, S. N., and D. Quah, “The New Empirics of Economic Growth,” in J. Taylor and M. Woodford, eds., *Handbook of Macroeconomics* (Amsterdam: North-Holland, 1999).
- Feenstra, R. C., R. Inklarr, and M. P. Timmer, “The Next Generation of the Penn World Tables,” *American Economic Review* 105:10 (2015), 3150–3182. 10.1257/aer.20130954
- Frühwirth-Schnatter, S. and Kaufmann, S., “Model-Based Clustering of Multiple Time Series,” *Journal of Business and Economic Statistics* 26:1 (2008), 78–89.
- Hamilton, J. D., and M. T. Owyang, “The Propagation of Regional Recessions,” *this REVIEW* 94:4 (2012), 935–947.
- Harvey, A. C., *Forecasting, Structural Time Series Models and the Kalman Filter* (Cambridge: Cambridge University Press, 1989).
- Hausmann, R., L. Pritchett, and D. Rodrik, “Growth Accelerations,” *Journal of Economic Growth* 10 (2005), 303–329. 10.1007/s10887-005-4712-0
- International Monetary Fund, *World Economics Outlook* (Washington, DC: International Monetary Fund, 2020).
- Johnson, P., and C. Papageorgiou, “What Remains of Cross-Country Convergence?” *Journal of Economic Literature* 58:1 (2020), 129–175. 10.1257/jel.20181207
- Jones, B. F., and B. A. Olken, “The Anatomy of Start-Stop Growth,” *this REVIEW* 90 (2008), 582–587.
- Jones, C. I., “On the Evolution of the World Income Distribution,” *Journal of Economic Perspectives* 11:3 (1997), 19–36. 10.1257/jep.11.3.19
- “The Facts of Economic Growth” (pp. 3–68), in J. B. Taylor and H. Uhlig, eds., *Handbook of Macroeconomics*, vol. 2A (Amsterdam: Elsevier, 2016). 10.1016/bs.hesmac.2016.03.002
- Kremer, M., A. Onatski, and J. H. Stock, “Searching for Prosperity,” *Carnegie Rochester Conference Series on Public Policy* 55 (2001). 10.1016/S0167-2231(01)00060-4
- Lucas, R. E., “Some Macroeconomics for the 21st Century,” *Journal of Economic Perspectives* 14:1 (2000), 159–168. 10.1257/jep.14.1.159
- Mankiw, N. G., D. Romer, and D. N. Weil (1992), “A Contribution to the Empirics of Economic Growth,” *Quarterly Journal of Economics*, 107:2 (1992), 407–437. 10.2307/2118477
- Moench, E., S. Ng, and S. Potter, “Dynamic Hierarchical Factor Models,” *this REVIEW* 95:5 (2013), 1811–1817.
- Müller, U. K., and M. W. Watson, “Testing Models of Low-Frequency Variability,” *Econometrica* 76:5 (2008), 979–1016.
- “Measuring Uncertainty about Long-Run Predictions,” *Review of Economic Studies* 83:4 (2016), 1711–1740.
- “Long-Run Covariability,” *Econometrica* 86:3 (2018), 775–804.
- “Low-Frequency Analysis of Economic Time Series,” unpublished manuscript, Department of Economics, Princeton University (2020).
- National Academy of Sciences, Committee on Assessing Approaches to Updating the Social Cost of Carbon, *Valuing Climate Changes: Updating Estimation of the Social Cost of Carbon Dioxide* (Washington, DC: National Academies Press, 2017), <https://www.nap.edu/24651>.
- Pritchett, L., “Understanding Patterns of Economic Growth: Searching for Hills among Plateaus, Mountains, and Plains,” *World Bank Economic Review* 14 (2000), 221–225. 10.1093/wber/14.2.221
- Quah, D., “Empirical Cross-Section Dynamics in Economic Growth,” *European Economic Review* 37 (1993), 426–434. 10.1016/0014-2921(93)90031-5

- “Twin Peaks: Growth and Convergence in Models of Distribution Dynamics,” *Economic Journal* 106 (1996), 1045–1055. 10.2307/2235377
- “Empirics for Growth and Distribution” Polarization, Stratification, and Convergence Clubs,” *Journal of Economic Growth* 2:1 (1997), 27–59. 10.1023/A:1009781613339
- Raftery, A. E., A. Zimmer, D. M. W. Frierson, R. Startz, and P. Liu, “Less than 2°C Warming by 2100 Unlikely,” *Nature: Climate Change* 7:99 (2017), 637–641. 10.1038/nclimate3352, PubMed: 30079118
- Ramey, G., and V. A. Ramey, “Cross-Country Evidence on the Link Between Volatility and Growth,” *American Economic Review* 85:1 (1995), 1138–1151.
- Sala-i-Martin, X., “I Just Ran Two Million Regressions,” *American Economic Review Papers and Proceedings* 82:2 (1997), 178–183.
- Startz, R., “The Next Hundred Years of Growth: Growth and Convergence,” *Journal of Applied Econometrics* 35 (2020). 10.1002/jae.2732
- Stock, J. H., and M. W. Watson, “Asymptotically Median Unbiased Estimation of Coefficient Variation in a Time Varying Parameter Model,” *Journal of the American Statistical Association* 93 (1998), 349–358. 10.1080/01621459.1998.10474116