Is Growth Additive?
A Bayesian Model Analysis and Evidence from Forecasts*

Callum Jones† David López-Salido‡ Thomas Philippon§

February 2026

Abstract

Growth theory is based on the assumption of exponential total factor productivity (TFP) growth, or that the expected size of the next TFP increment is proportional to the current level of TFP. Across countries and time periods we find strong evidence that TFP growth is additive and not exponential. Even starting from low priors, Bayesian estimation selects the additive model over the exponential one. Moreover, we show that professional forecasts are in line with the additive model, not the exponential one, as the additive growth model, unlike the exponential one, provides useful long-term forecasts for TFP.

Keywords: growth theory, TFP, additive growth.

*We are grateful to Tim Cogley, Chad Jones, Ufuk Akcigit, Virgiliu Midrigan, Greg Mankiw, Olivier Blanchard, Xavier Gabaix, Xavier Jaravel, Antonin Bergeaud, Remy Lecat, David Weil, Gilbert Cette, David Romer, German Gutierrez, Ben Jones, Alexey Guzey, Peter Kruse-Andersen, and Bill Easterly for their comments, and to Yad Selvakumar and Nicholas Zevanove for outstanding research assistance. The views expressed are those of the authors and not necessarily those of the Federal Reserve Board or the Banco de España.

†Federal Reserve Board, callum.j.jones@frb.gov

‡Banco de España and CEPR, david.lopez-salido@bde.es

§New York University Stern School of Business, CEPR and NBER, tphilipp@stern.nyu.edu

Introduction

We propose an empirical investigation of the stochastic process that governs total factor productivity (TFP). At least since Solow (1956) economists have assumed that TFP follows an exponential process (henceforth, model $\mathcal{G}$ for "geometric") which takes the form:

$$\log A_{t+\tau} = \log A_t + g_t \tau + noise, \tag{1}$$

where A_t is TFP in year t and the trend component g_t is constant or at least highly persistent. We will show instead that growth is additive and that the TFP process is better described by model $\mathcal{A}$ (as in "additive" or "arithmetic"):

$$A_{t+\tau} = A_t + b_t \tau + noise, \tag{2}$$

where b_t is constant or at least highly persistent. The key prediction of model $\mathcal{G}$ is that the expected size of the *next* TFP increment is proportional to the *current level* of TFP. We examine data across many countries and time periods and find that this prediction is rejected. In most cases productivity growth appears to be additive.

Figure 1 provides an informal but transparent motivation for our paper. It suggests that TFP growth has been nearly linear in the US since at least World War II. In the figure, models $\mathcal{A}$ and $\mathcal{G}$ are estimated over the first half of the sample (1947-1983) and then used to predict the level of TFP in the second half of the sample (1984-2019). We observe the well-known TFP slowdown "puzzle" with model $\mathcal{G}$, which simply says that actual TFP has fallen short of the exponential benchmark. By contrast there is no TFP slowdown according to model $\mathcal{A}$. To see this, equivalently, Figure 2 shows that the additive growth model predicts the correct path of TFP slowdown.

Empirically, then, US growth after World War 2 is well described by the following statement: Hicks-neutral TFP, normalized to 1 in 1947, increases each year by about 250 basis points. The *initial* trend growth rate is 2.5% but growth is additive: as TFP doubles after 40 year and increments are constant, the measured trend growth rate in percent is half of what it used to be. After 60 years, it is around one percent, in line with the data. When Hicks-neutral TFP grows linearly, capital accumulation creates a convex path for labor productivity. The Appendix shows that the linear TFP model predicts the correct non-linear evolution of labor productivity while the exponential model over-predicts future levels of labor productivity, just like it does for TFP. The point that TFP is additive therefore does not depend on a (possibly noisy) measure of the capital stock.

The goal of our paper, then, is to formally test the proposition that TFP growth is

Figure 1: US Post War TFP

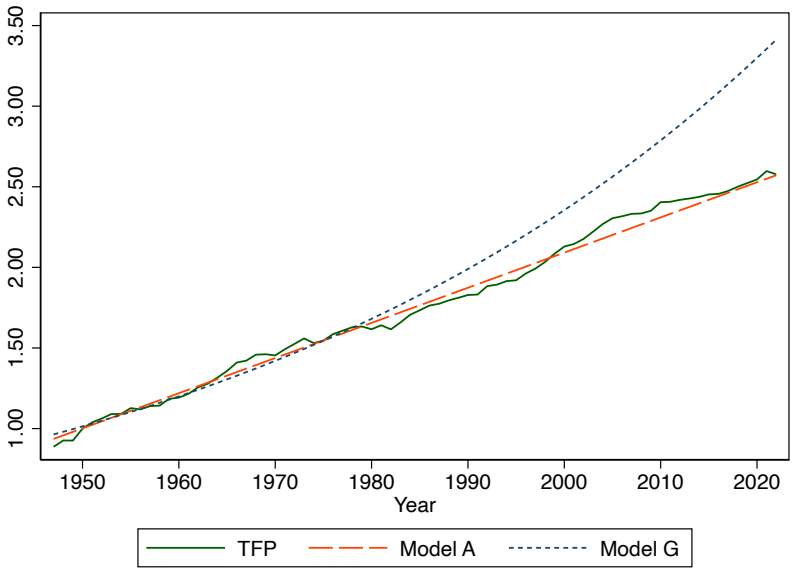

Notes: TFP is normalized to 1 in 1947. Models are estimated over 1947-1983. The forecast 1984-2019 is out-of-sample. Data source: Bergeaud et al. (2016). described in Appendix A.

additive, across many countries and time periods. The challenge, as clearly seen in equations (1) and (2), is to specify the "noise" and obtain estimates of the stochastic trends g_t and b_t. Section 1 presents our statistical framework.

Section 2 focuses on US data and estimates the models $\mathcal{A}$ and $\mathcal{G}$ using maximum likelihood techniques. We then filter the unobserved states g_t and b_t with a Kalman filter, before approximating the conditional expectations using Monte Carlo simulations. The estimates reveal that the trend in Figure 1 actually starts around 1930. This finding is consistent with the historical literature (Field, 2003; David, 1990; Gordon, 2016), and it suggests that US TFP growth has been additive with approximately constant expected increments for 90 years. Using these estimates, Section 3 conducts a Bayesian model comparison exercise. The exercise confirms the visual impression of Figure 1.

Model $\mathcal{G}$ performs poorly even when the trend g is allowed to change over time, either continuously or with discrete breaks. Thus, in the model selection exercise, the posterior model probability converges to one for model $\mathcal{A}$ by 2022, the end of our sample, even when one starts with large priors in favor of model $\mathcal{G}$ in 1890. In section 6, we report that these conclusions hold across different estimation approaches as well as alternative specifications of models $\mathcal{A}$ and $\mathcal{G}$. In particular, we show that the posterior weight of model $\mathcal{A}$ still converges to one when the trend growth process is autoregressive as opposed to a random walk, and also when we allow the trend growth process to instead be subject to regime breaks at

Figure 2: Predicted TFP Slowdown

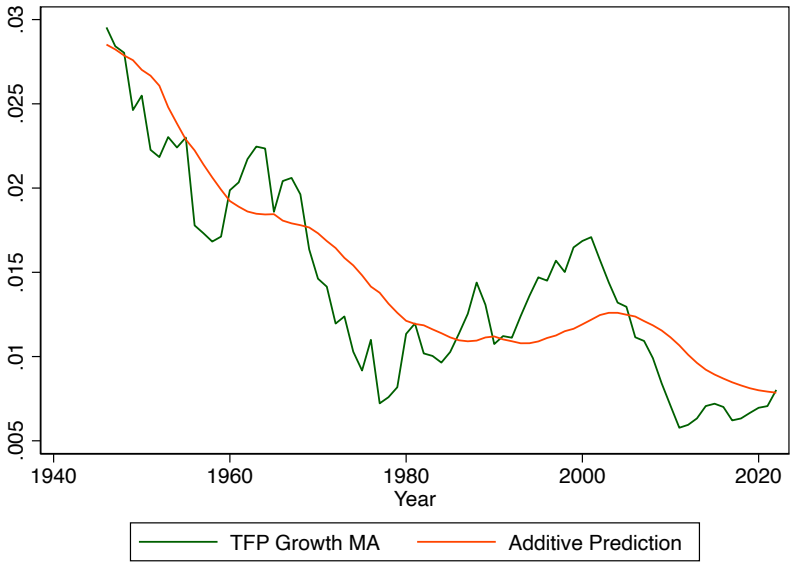

Notes: TFP is normalized to 1 in 1947. TFP Growth MA is the centered moving average of TFP growth over $(t-5, t+5)$. The prediction of model $\mathcal{A}$ is based on a constant annual increment of 250 basis points. Data source: Bergeaud et al. (2016).

estimated periods. We also show that the results hold when we use a Bayesian estimation and an estimation with a recursively expanding sample.

Economists often favor the geometric model for its theoretical simplicity rather than its accuracy, while professional forecasters choose models they find most plausible, aiming for predictions aligned with actual growth patterns. Drawing on Milton Friedman billiard player analogy, this confidence stems not from assuming forecasters compute precise Bayes factors, but from the expectation that their expertise ensures results close to reality. To test this, in Section 4 we use labor productivity forecasts from the Council of Economic Advisors (CEA) and Survey of Professional Forecasters (SPF), convert them to TFP forecasts under a Cobb-Douglas framework, and evaluate their log predictive scores, cumulative Bayes factors, and posterior probabilities to determine which model best explains these forecasts.

Our findings show that model $\mathcal{A}$ aligns well with SPF and CEA forecasts, except during 2017-2021 when CEA projections were overly optimistic. Interestingly, model $\mathcal{A}$ not only mirrors professional forecasts but also outperforms them in predicting actual TFP, particularly between 2010 and 2020. Bayesian inference reinforces this conclusion: professional forecasters overwhelmingly rely on a model similar to $\mathcal{A}$, with probabilities near one. For CEA, this probability also approaches one but declines after the optimistic 2017-2021 forecasts, suggesting limited convexity in their TFP expectations.

Section 5 repeats the analysis of Sections 2 and 3 in the panel of 23 countries from

Bergeaud et al. (2016). The additive model predicts TFP dynamics better than the exponential model for most of the 23 countries. The 10-year forecast errors of the exponential model are 30% to 60% higher than those of the additive model. We also consider a sample of Asian countries that are not in the BCL sample (e.g., Korea, Taiwan, and China). We find no evidence of sustained convexity in TFP, even during periods commonly described as exhibiting "miracle growth", and strong evidence of additive growth (e.g., Taiwan).

Literature Our paper uses a Bayesian framework to derive time-varying weights on candidate growth models that adjust as new data arrives. The approach relates to that in Del Negro et al. (2016), which tackles model misspecification by recognizing that no single model perfectly captures economic dynamics; instead, it distributes credibility across models and shifts weights when evidence suggests changes in predictive performance. By updating weights through model forecasting performance, the approach uses Bayesian inference to manage uncertainty and improve robustness.

Our paper also relates with Edge et al. (2007) that examine how private-sector forecasts influence perceptions of long-run U.S. productivity growth and its macroeconomic implications. In the context of a standard geometric growth model, they highlight that shifts in trend productivity—such as those in the 1970s and 1990s—are hard to detect in real time, often confounded by cyclical noise. Our additive growth framework reinforces the idea that private sector long run forecasts provides valuable information—in addition to actual data—to identify the presence of additive growth. Our paper highlights the importance of private sector expectations to better match historical data and avoids overestimating future medium-to-longer run productivity growth.

Solow (1956) studies the theoretical properties of the neoclassical growth model and Solow (1957) constructs TFP series from 1909 to 1949. Since then, essentially all models of growth have taken the exponential model as a benchmark. Our Bayesian estimation rejects the exponential model using the full sample, but, interestingly, the power of the test is lower if we only use Solow's original data because of the structural break in 1930.

This paper complements the literature on endogenous growth accounting, such as Jones (2002). Compared to Solow (1957), this literature treats TFP growth as an endogenous variable to be explained by inputs such as capital, education, and the labor force employed in research. A key puzzle in the literature is that TFP growth has not increased despite the increase in measured research effort (Jones, 1995). Jones (2009) argues that innovation is getting harder because new generations of innovators face an increasing educational burden, while Akcigit and Ates (2023) argue that increasing barriers to knowledge diffusion across firms can explain the evolution of business entry, exit and income shares. Bloom et al. (2020)

present case studies of several technologies to argue that innovations are becoming harder to find. Guzey et al. (2021), however, show that this conclusion is sensitive to the choice of a productivity measure, and that many series, including US TFP, do not appear to exhibit exponential growth.

A critical issue in this literature is the measurement of inputs into the innovation process. Formal R&D spending captures only a fraction of innovative activities, and R&D data is usually missing before World War 2. Growth accounting also depends on the functional form chosen to map inputs into TFP. By contrast, our approach focuses directly on measured TFP and shows that its stochastic process is additive across most countries and time periods. This new stylized fact speaks to all models of growth, endogenous or semi-endogenous, based on R&D or on learning-by-doing.

This paper is not the first to suggest a departure from exponential growth. Jones (1995), for instance, includes a TFP equation of the type $\dot{A}_t = A_t^{\phi} L_{A,t}$ where $L_{A,t}$ is research employment. Exponential growth in standard models comes from the assumption that $\phi = 1$. The endogenous growth accounting literature calibrates $\phi < 1$ to match the fact that increasing research effort does not necessarily lead to faster growth, but the estimates of ϕ using R&D data are rather unstable. Our results suggest that ϕ is close to zero in the aggregate (not necessarily at the firm level as we discuss in conclusion).

Finally this paper relates to the history of long run growth. The fact that growth increments increase during industrial revolutions speaks to the complementarity of new inventions with existing technologies emphasized by Comin et al. (2010). The turning point of the 1930s is consistent with Field (2003)'s argument that "*the years 1929–1941 were, in the aggregate, the most technologically progressive of any comparable period in U.S. economic history.*"

1 Modeling TFP Growth

The challenge is to create a test that can distinguish between additive and geometric growth. Let us start with an informal example to build some intuition. Suppose that measured TFP is equal to fundamental TFP A_t^* plus noise $A_t = A_t^* + \epsilon_t$ where ϵ_t is *iid* with volatility σ_ϵ. Let us normalize $A_0^* = 1$, assume that fundamental TFP is deterministic and that we know the trend g at time 0. We do not know whether A_t^* will grow exponentially as $A_t^* = e^{gt}$, or linearly as $A_t^* = 1 + gt$. How long would we need to wait before deciding which is the correct model? If we want the two forecasts to be $n\sigma$ apart we need $e^{gt} - 1 - gt \geq n\sigma$ or $gt \geq x$ where x is the root of $e^x - 1 - x = n\sigma$. For small values of $n\sigma$, x is close to $\sqrt{2n\sigma}$ so we need $t \geq \frac{\sqrt{2n\sigma}}{g}$. With $n = 2$, $\sigma = 1\%$ and $g = 2\%$ we get approximately $t \geq 10$ years. If we assume instead that the noise is multiplicative $A_t = A_t^* (1 + \epsilon_t)$ then we need approximately $t \geq 12$

years. This back-of-the-envelope calculation gives a sense of the "useful" forecast horizon to test convexity of the TFP growth process. Let us now proceed with a formal Bayesian analysis.

1.1 Benchmark Specification

We consider two benchmark models for the evolution of observed TFP A_t. We specify model $\mathcal{A}$ (additive growth) as

$$\begin{aligned} A_t &= A_{t-1} + b_t + \sigma_a \epsilon_t^a \\ b_t &= b_{t-1} + \sigma_u u_t \end{aligned}$$

where u_t and ϵ_t^a are independent standard normal shocks. We specify model $\mathcal{G}$ (geometric growth) as

$$\begin{aligned} A_t &= A_{t-1}\left(1 + g_t\right) \exp\left(\sigma_g \epsilon_t^g\right) \\ g_t &= g_{t-1} + \sigma_\nu \nu_t \end{aligned}$$

where ν_t and ϵ_t^g are independent standard normal shocks. In particular $\log A_t$ is additive, as in equation (1). In the Robustness section below we consider various extensions for the U.S. data such as treating b_t and g_t instead as parameters and allowing for regime shifts in those parameters.

1.2 State space representation

We write each model $m = \mathcal{A}, \mathcal{G}$ in standard VAR(1) representation as

$$x_t^m = \mathbf{Q}_m x_{t-1}^m + \mathbf{G}_m \epsilon_t^m. \tag{3}$$

For $m = \mathcal{A}$ we have $A_t = A_{t-1} + b_{t-1} + \sigma_u u_t + \sigma_a \epsilon_t^a$, so that

$$x_t^{\mathcal{A}} \equiv \begin{bmatrix} A_t \\ b_t \end{bmatrix}, \quad \epsilon_t^{\mathcal{A}} \equiv \begin{bmatrix} \epsilon_t^a \\ u_t \end{bmatrix}, \quad \mathbf{Q}_{\mathcal{A}} \equiv \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{G}_{\mathcal{A}} \equiv \begin{bmatrix} \sigma_a & \sigma_u \\ 0 & \sigma_u \end{bmatrix}.$$

For $m = \mathcal{G}$ we have $\log A_t = \log A_{t-1} + g_{t-1} + \sigma_\nu \nu_t + \sigma_g \epsilon_t^g$ and $g_t = g_{t-1} + \sigma_\nu \nu_t$, so that

$$x_t^{\mathcal{G}} = \begin{bmatrix} \log A_t \\ g_t \end{bmatrix}, \quad \epsilon_t^{\mathcal{G}} = \begin{bmatrix} \epsilon_t^g \\ \nu_t \end{bmatrix}, \quad \mathbf{Q}_{\mathcal{G}} \equiv \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{G}_{\mathcal{G}} \equiv \begin{bmatrix} \sigma_g & \sigma_\nu \\ 0 & \sigma_\nu \end{bmatrix}.$$

Figure 3: Estimation Sample, U.S. TFP

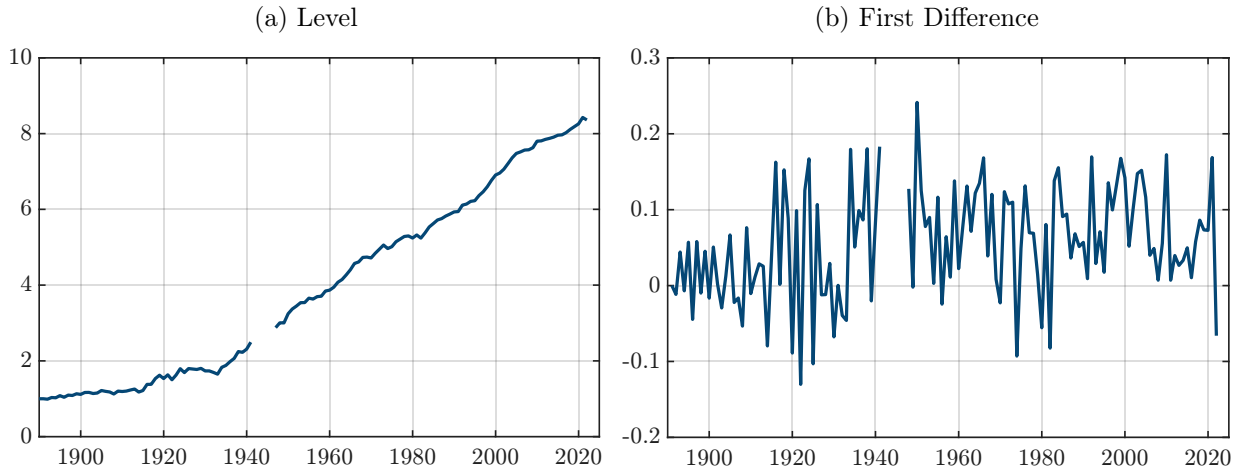

Notes: The left panel shows the level of TFP from BCL, which is normalized to 1 in 1890. The right panel shows the first difference in TFP. The WW2 years 1942-1946 are omitted from the sample.

With the two models written in their state space forms, we can then use standard maximum likelihood techniques to estimate the parameters of the two models.[1] We do this in the next section.

2 Estimation for the U.S.

2.1 Data

Data on TFP for the U.S. from Bergeaud et al. (2016) is available from 1890 to 2022. We omit the WW2 years 1942–1946. Appendix A describes the U.S. data in more detail. We add a lower barrier of zero for the stochastic trends b_t and g_t based on theoretical priors (i.e., no fundamental technological regress), and we normalize the value of TFP in the first year to one. Figure 3 plots the data we use in both levels and first differences.

2.2 Full Sample Estimation of Second Moments

As a benchmark, we estimate the parameters $\theta^{\mathcal{A}} = (\sigma_a, \sigma_u)$ for the additive model and $\theta^{\mathcal{G}} = (\sigma_g, \sigma_\nu)$ for the geometric model using full-sample maximum likelihood estimation (MLE). Note, however, that we only estimate the second moments of the data in the full sample, so the risk of creating look-ahead bias is small. Nonetheless, as a robustness, we

[1] We explore alternative ways to estimate the model parameters in Appendix B, including standard Bayesian techniques with uniform priors on the parameters, as well as the specification where we treat b_t and g_t as estimated parameters subject to regime shifts in periods that are also estimated.

Table 1: MLE Parameter Estimates (US, 1880-2022)

	Model $\mathcal{A}$		Model $\mathcal{G}$	
	σ_u	σ_a	σ_ν	σ_g
MLE Estimate	0.005	0.069	0.001	0.032
Standard Deviation	(0.002)	(0.005)	(0.001)	(0.002)

Notes: This table provides the parameter estimates for the parameters governing the standard deviations of the $\mathcal{A}$ and $\mathcal{G}$ models under MLE.

also run a real time MLE estimation over expanding samples and show that the results of our model evaluation exercise are quantitatively similar (we report these results in Appendix B). Table 1 shows the MLE estimates in our benchmark specification.

As an additional robustness exercise, we characterize the posterior distributions of the parameters $\theta^{\mathcal{A}}, \theta^{\mathcal{G}}$. To do so we estimate the σ's using Bayesian methods, i.e., a Markov Chain Monte Carlo (MCMC) procedure. We use uniform priors for each parameter. In our robustness exercises reported in Section 6 below, we provide details on the MCMC approach along with estimates of the posterior distributions, whose modal values and standard deviations, reported in Table 2, are close to the MLE estimates of Table 1.

2.3 Recursive Estimation of Stochastic Trends and Forecasts

We use a recursive, real-time approach to estimate b_t and g_t. Given the MLE parameter estimates above, for each forecast origin t, we filter the states using data $A_1, \ldots, A_t$ only. Figure 4 shows the real-time filtered estimates of b_t and g_t. The value of b_t is around 0.02 during the first part of the sample up until about the 1930s, after which it is estimated to increase to around 0.07 by 1950 and fluctuate around that value. The value of g_t is estimated to be around 0.015 until about the 1930s, rises to around 0.02 when growth accelerates after the 1930s, and then is estimated to trend down from the 1970s.

Using the estimates of the parameters and filtered states, we then create TFP forecasts $\{A_{t+h}\}_{h=1}^{30}$. Finally we use the forecasts to sequentially assess and compare the predictive performance of the models.

3 Bayesian Model Comparison

This section discusses the statistical approach we take in comparing the two models.

Figure 4: MLE Filtered States

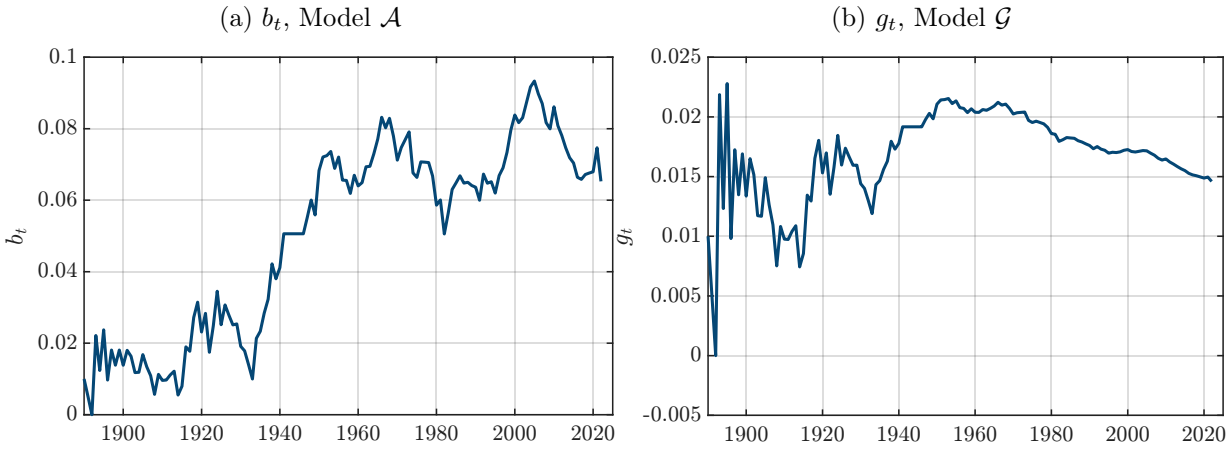

Notes: This figure shows the filtered b_t and g_t under the MLE parameter estimates.

3.1 Cumulative Log Predictive Score

We now use the benchmark MLE estimates presented above to compare the models. Appendix B shows that our results are robust under alternative estimation procedures. At each forecast origin we compute the predicted density of models $m = \mathcal{A}, \mathcal{G}$:

$$p\left(x \mid A^t, m, h\right) \equiv \Pr\left(A_{t+h} = x \mid A^t, m\right). \tag{4}$$

For both models $\mathcal{A}$ and $\mathcal{G}$, we use Monte Carlo simulations to generate the predictive densities. Figure 5 displays the results.

Figure 5 is the key figure of the paper. The black line is the actual path of US TFP. The blue lines are the median model forecasts and the shaded areas are the confidence intervals. We will discuss the comparison with the SPF and CEA forecasts in the next section. Of course, neither model can forecast the structural change in trend growth in the 1930s, so both models provide pessimistic forecasts of TFP in the 1950s. As the models filter the early post-war data they revise upward their estimates of b_t and g_t. But the medium / long term forecasts of model $\mathcal{G}$ are widely optimistic, in addition to being very imprecise. Model $\mathcal{A}$, on the other hand, is broadly consistent with the historical evidence. We now formalize and quantify these points.

We evaluate the relative performance of models $\mathcal{A}$ and $\mathcal{G}$ by comparing their out-of-sample forecast performance. Given *realized* productivity A_{t+h} the log predictive score is

$$\log p\left(A_{t+h} \mid A^t, m, h\right),$$

Figure 5: TFP Level and Forecasts under MLE Estimates

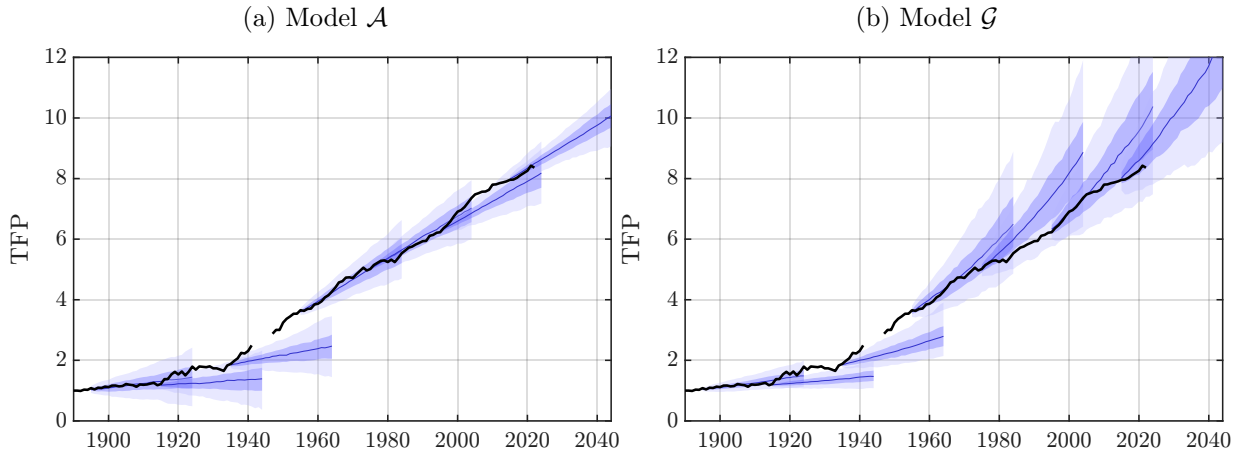

Notes: The blue line plots the model predicted median paths of TFP (every 20 years, over a forecast horizon of 30 years), along with the 5% / 95% range (in light blue) and the 25% / 75% range (in darker blue).

where p is defined in (4). The key statistic used is the *cumulative log score* (CLS), computed as the sum across time of the log predictive scores for h-step-ahead forecasts:

$$\text{CLS}_m(T) = \sum_{t=1}^{T-h} \log p\left(A_{t+h} \mid A^t, m, h\right). \tag{5}$$

We use $h = 10$ years as our benchmark forecast horizon and therefore T runs from 1900 to 2022. Our results are robust to using longer horizons, as shown in Appendix B for $h = 15$ years.

3.2 Log Bayes Factor and Posterior Model Probabilities

From the cumulative log scores in (5), the log Bayes factor in favor of Model $\mathcal{A}$ is

$$\log B_{\mathcal{A},\mathcal{G}}(T) = \text{CLS}_{\mathcal{A}}(T) - \text{CLS}_{\mathcal{G}}(T). \tag{6}$$

When $\log B_{\mathcal{A},\mathcal{G}}(T) > 0$, Model $\mathcal{A}$ has superior cumulative forecast accuracy; when negative, Model $\mathcal{G}$ does.

Figure 6, Panel (a) shows the log Bayes factors in favor of model $\mathcal{A}$. Panel (a) shows that the data up to 1960 is inconclusive for two main reasons. Firstly, TFP growth before WWI is rather slow so the convexity effect is relatively small. Second, as is well known from historical research (Field, 2003; David, 1990; Gordon, 2016) there is a structural shift in the 1930s with the large scale implementation of several new technologies. The positive change

in the slope of the TFP series makes it difficult to differentiate the two models. By the late 1960s, however, the lack of convexity in the TFP series starts to become more obvious to our Bayesian estimator.

Panel (b) of Figure 6 translates the log Bayes factors into posterior probabilities, using an adjustment proposed by Del Negro et al. (2016). Let us start with standard Bayesian Model Averaging (BMA). If we assign prior probabilities $\pi_{\mathcal{A}}$ and $\pi_{\mathcal{G}} = 1 - \pi_A$ to the two models, then the posterior probability of Model $\mathcal{A}$ given sample T using the standard BMA formula is

$$\mathbb{P}_T(\mathcal{A}) = \frac{\pi_A \cdot B_{\mathcal{A},\mathcal{G}}(T)}{\pi_A \cdot B_{\mathcal{A},\mathcal{G}}(T) + \pi_G} = \frac{\pi_A \cdot \exp\left[\mathrm{CLS}_{\mathcal{A}}(T) - \mathrm{CLS}_{\mathcal{G}}(T)\right]}{\pi_A \cdot \exp\left[\mathrm{CLS}_{\mathcal{A}}(T) - \mathrm{CLS}_{\mathcal{G}}(T)\right] + \pi_G}. \tag{7}$$

As Del Negro et al. (2016) point out, however, standard BMA can produce unreasonable swings in posterior probabilities. There are two broad classes of explanations for this problem. First, BMA assumes that exactly one of the two models is correct, which is not plausible in most economic studies, ours included. At most we wish to claim that the linear process provides a useful approximation of the true TFP process and that we can reject the strong convexity of the geometric process, but obviously there are models with weaker convexity that we could not reject. Second, the estimation, including Kalman filtering, assumes normal disturbances. The filtering itself is robust to non-normality since it is an orthogonal projection in a Hilbert space, but the implied posterior probabilities are sensitive to the specification of the error terms and normality gives a false sense of precision: a CLS of ± 3 can easily happen with noisy data, but since $\exp(3) \approx 20$, one would conclude strongly in favor of one model when no such confidence is in fact warranted.

Del Negro et al. (2016) propose instead a dynamic prediction pool and a scaling factor γ that tempers the elasticity of the posterior probability to the log Bayes factor. We use the following specification for the posterior probability weight in the dynamic prediction pool:

$$w_T = \frac{\pi_A \cdot \exp\left[\gamma\left\{\mathrm{CLS}_{\mathcal{A}}(T) - \mathrm{CLS}_{\mathcal{G}}(T)\right\}\right]}{\pi_A \cdot \exp\left[\gamma\left\{\mathrm{CLS}_{\mathcal{A}}(T) - \mathrm{CLS}_{\mathcal{G}}(T)\right\}\right] + \pi_G}. \tag{8}$$

A small value of γ makes the posterior weight react less sharply to recent differences in forecast accuracy. In our exercises below, we set γ to a conservative value of 0.25. Under this value, simulation exercises show that it takes about 20 years for cumulative log score differences of around 1 each period to raise a posterior model weight from 0.1 to near 1. In the estimation below the main effect of $\gamma = 0.25$ is to lower the volatility of the estimated posterior weight in the 1930s. Clearly the conclusion that the weight ≈ 1 when $T > 1980$ would only be stronger with $\gamma = 1$.

Figure 6: Log Bayes Factor and Posterior Model Weight

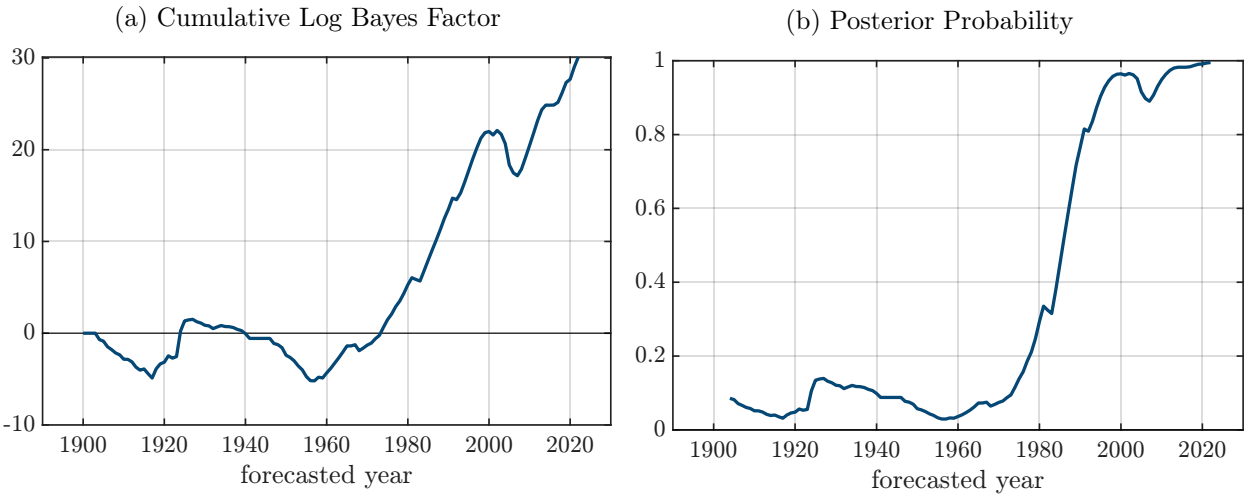

Notes: When computing the posterior probability, the prior of model $\mathcal{G}$ is 0.9. The vertical axes show the cumulative log Bayes factor and the posterior model weight at the date t on the horizontal axis for a forecast made at $t-10$. We use $\gamma = 0.25$ in constructing the posterior probability.

As shown in panel (b) of Figure 6, the posterior weight of model $\mathcal{A}$ under (8) hovers around the prior of 0.1 until about the 1970s. From then on, model $\mathcal{A}$ shows superior forecast accuracy for future TFP, with the cumulative log Bayes factor shifting strongly towards it. Apart from a short dip in the 2000s, the posterior weight of model $\mathcal{A}$ is near 1.[2]

3.3 Revisiting Solow (1957)

> "Not only is $\Delta A/A$ uncorrelated with K/L, but one might almost conclude from the graph that $\Delta A/A$ is essentially constant in time, exhibiting more or less random fluctuations about a fixed mean."
>
> Solow (1957)

The Bayesian approach to model selection can also shed light on the history of economic research on growth. Using data from 1909 to 1949, Solow (1957) found a pattern for A_t qualitatively similar to that in Figure 4. He wrote that "*there does seem to be a break at about 1930. There is some evidence that the average rate of progress in the years 1909-29 was smaller than that from 1930-49.*" Indeed, a formal test finds a structural break around 1930 in the first difference series $A_t - A_{t-1}$. The change in the slope makes it difficult to distinguish models $\mathcal{A}$ and $\mathcal{G}$ using only Solow's data. Formally, if one feeds data from 1909

[2]In Appendix B.1, we complement this density-based evaluation with a point-forecast comparison based on mean squared error. This simpler exercise abstracts from predictive uncertainty. We find that the additive specification continues to outperform the geometric model in MSE terms.

to 1949, the posteriors over the two models are not very different from the priors. Thus Solow made a reasonable choice based on available evidence at the time of his writing, and it is only with the benefit of post-1950 data that we can confidently conclude that growth is additive.

4 Comparison with Professional Forecasts

> "Our confidence in this hypothesis is not based on the belief that billiard players, even expert ones, can or do go through the process described; it derives rather from the belief that, unless in some way or other they were capable of reaching essentially the same result, they would not in fact be expert billiard players."
>
> Milton Friedman, Essays in Positive Economic, 1953

Economists have used the geometric model for its theoretical convenience, not for its accuracy. Professional forecasters, on the other hand, have incentives to use whichever model they find most plausible. In doing so, we might expect them to produce forecasts that are most consistent with the actual process of growth. To borrow from Milton Friedman's expert billiard player metaphor, our confidence in this hypothesis is not based on the belief that professional forecasters have derived the correct cumulative Bayes factors, but rather that, if they were not able to reach approximately the same result, they would not be professional forecasters.

We thus use forecasts of labor productivity from the CEA and the SPF to test our model of growth (see the data Appendix A for details of each forecast). We convert them to TFP forecasts assuming a Cobb-Douglas production function. We then compute the log predictive score of the forecast, under the full sample MLE estimates of each model, and evaluate the corresponding cumulative log Bayes factor and posterior model probability. This provides a measure of which model better explains the predictions made by professional and CEA forecasters.

Figure 7 shows our results. Panel (a) presents the data on CEA and SPF forecasts, alongside the range of estimates of TFP from model $\mathcal{A}$, with *forecasted* year on the time axis. We also plot realized TFP in those years in panel (a). SPF forecasts are available from 1992:Q1 to 2022:Q3 at a consistent 10 year horizon. The first star, in 2002, corresponds to the forecast made in 1992. CEA forecasts are available as early as 1970 but with inconsistent time horizons, as we explain in Appendix A.

Panel (a) shows that model $\mathcal{A}$ is broadly consistent with SPF forecasts and with CEA forecasts except for the 2017-2021 period where CEA forecasts appear optimistic in a way

Figure 7: Comparison with Forecasters

(a) Observed TFP, and Forecasts $(t-10) \to t$ from: Model $\mathcal{A}$, SPF, and CEA

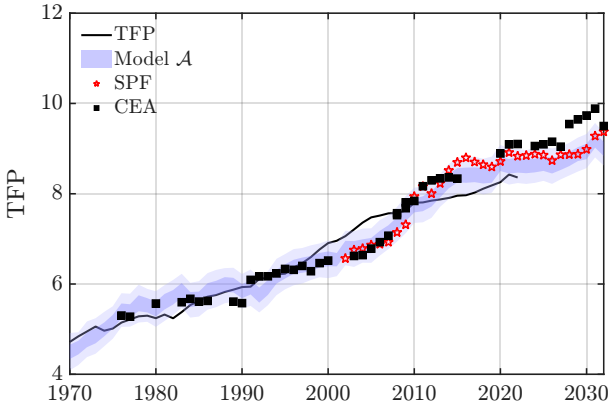

(b) Posterior Probability of Model $\mathcal{A}$, SPF and CEA Forecast Evaluation under MLE Estimates

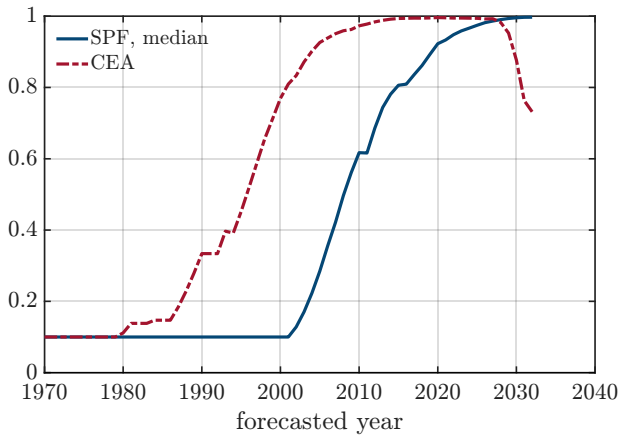

Notes: Year t is the forecasted year. The shaded areas are the 25% and 75% bands (solid) and 5% and 95% bands (lighter shade) for model $\mathcal{A}$ based on data up to $t-10$. Stars and Squares are median SPF and CEA forecasts for year t. Note that growth forecasts for the CEA are for different horizons each year (i.e., sometimes less than 10 years), as described in the data Appendix A. In panel (b), the initial prior on the geometric model is 0.9.

which is inconsistent both with model $\mathcal{A}$ and with SPF forecasts. Thus it appears that actual forecasters do not build in much convexity in their TFP forecasts. Second, and more surprisingly, model $\mathcal{A}$ provides better forecasts of actual TFP than either SPF and CEA: whenever there is a discrepancy between model A and SPF, model A is actually closer to the true realized productivity. This seems especially true between 2010 and 2020.

Panel (b) translates the forecasts and CLS into posterior probabilities. As with the actual data, Bayesian inference concludes that professional forecasters use a model similar to model $\mathcal{A}$ with probability close to one. For CEA forecasts, the probability converges to almost one but decreases with the release of the 2017-2021 forecasts, as CEA economists publish unexpectedly optimistic forecasts of TFP during those years.

5 International Evidence

Bergeaud et al. (2016) provide data for 23 countries.[3] We now apply the models of Section 2 to each country. This exercise is useful for two reasons: first, the sample is larger; second, we

[3]Australia, Austria, Belgium, Canada, Switzerland, Chile, Germany, Denmark, Spain, Finland, France, United Kingdom, Greece, Ireland, Italy, Japan, Mexico, Netherlands, Norway, New Zealand, Portugal, Sweden, and the United States. We also use provided data for a Euro Area aggregate. The sample covers the post-WW2 period 1950–2022. Appendix C plots the TFP paths for the countries used in the analysis.

Figure 8: Model Comparison, BCL Dataset

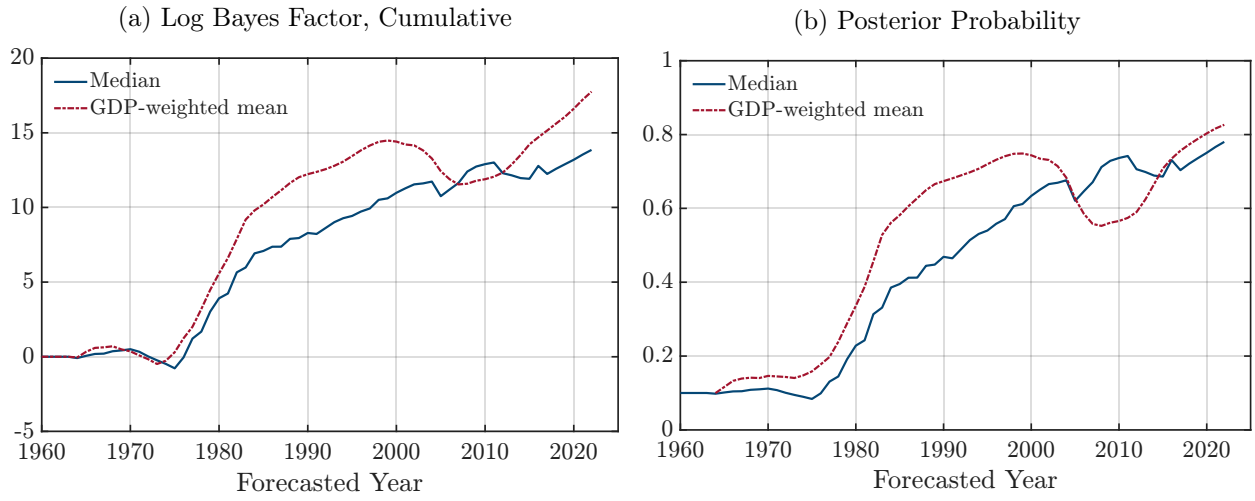

Notes: The line plots the median across countries in the BCL dataset and the real GDP weighted mean across countries. The initial prior on the geometric model is 0.9.

can investigate TFP growth in countries that are not necessarily at the frontier of technology in all industries. We find that model $\mathcal{A}$ beats model $\mathcal{G}$ in almost all countries, and often by a wider margin that in the US.

Figure 8 summarizes the international evidence from the BCL dataset by plotting the average across countries of the cumulative log Bayes factors and posterior model probabilities. The median log Bayes factor becomes positive by the late 1970s, indicating that, on average, the additive model delivers superior 10-year-ahead forecasts across countries. In addition, most countries in the sample show a positive log Bayes factor, with the notable exception of Ireland where GDP-based measures differ significantly from GNI-based measures.[4]

We next extend the analysis to a selection of fast-growing Asian economies using Penn World Table data: Korea, Taiwan, and China. Figure 9 shows that, despite very different initial conditions and episodes of rapid catch-up, TFP paths in Korea and Taiwan are approximately linear in levels over extended periods. In particular, there is little evidence of sustained convexity in TFP, even during phases commonly described as "miracle growth."

Figure 10 quantifies these visual impressions using posterior probabilities for forecast horizons of 10 and 20 years given the model estimates. We use as the initial prior for the geometric model in each country the posterior probability estimated for the U.S. at the date that the forecast is made.[5] For Korea and Taiwan, the cumulative log Bayes factor

[4]We make no attempt to correct the data for well-known issues, such as GNI vs GDP in Ireland, to avoid any risk of bias. We simply use the GDP-based TFP estimated from Bergeaud et al. (2016). Individual country results are shown in Appendix C.

[5]Under this, the initial prior probability for the geometric model is around 0.9, which was our assigned

rises steadily, implying support for the additive growth model at both horizons. Taiwan is probably the most striking example of linear TFP growth. The evidence for China is – unsurprisingly – ambiguous. China experiences a structural break around 1980, from stagnation to strongly increasing TFP, and there is some growth acceleration between 2000 and 2010. In that sense Chinese TFP around 1980 is similar to US TFP around 1930, and the Bayesian model does not draw a strong conclusion.

6 Extensions and Robustness

In this section, we consider a number of extensions of the baseline framework using U.S. data. We show that our results are robust to estimating the model in real-time, to a Bayesian estimation of the parameters, and to alternative specifications, including an autoregressive assumption for the trend growth terms as well as treating the trend growth terms as parameters subject to regime breaks.

6.1 Recursive MLE Estimation

The baseline results used the full sample MLE parameter estimates to evaluate the models. In this section, we instead evaluate the models using recursive estimates based on expanding window. That is, each period we expand the dataset by one year, recompute the MLE estimates of the model parameters on this dataset, and construct the log scores and cumulative log Bayes factors using the period-by-period estimates. Figure 11 shows the results evaluated at a 10-year forecast horizon, with our baseline full sample estimates also plotted for comparison. As the figure shows, by the end of the sample, the $\mathcal{A}$ model is preferred with a posterior model weight ≈ 1, as was the case under the full sample estimates. The recursive estimation also reinforces the point made in section 3.3 about the observations made by Solow (1957) about $\Delta A/A$ being roughly constant over time, with the strongest evidence for geometric growth occurring around 1950.

6.2 Bayesian Estimation

Our baseline procedure used MLE to estimate the model parameters. In this section, we instead do a Bayesian estimation of the $\mathcal{A}$ and $\mathcal{G}$ model parameters using the U.S. data. One benefit of conducting this estimation is that it yields posterior distributions for the model parameters with which we can conduct the Bayesian model analyses as above in the baseline.

prior on the geometric model in the exercises above.

Figure 9: TFP, Select Countries in PWT Dataset

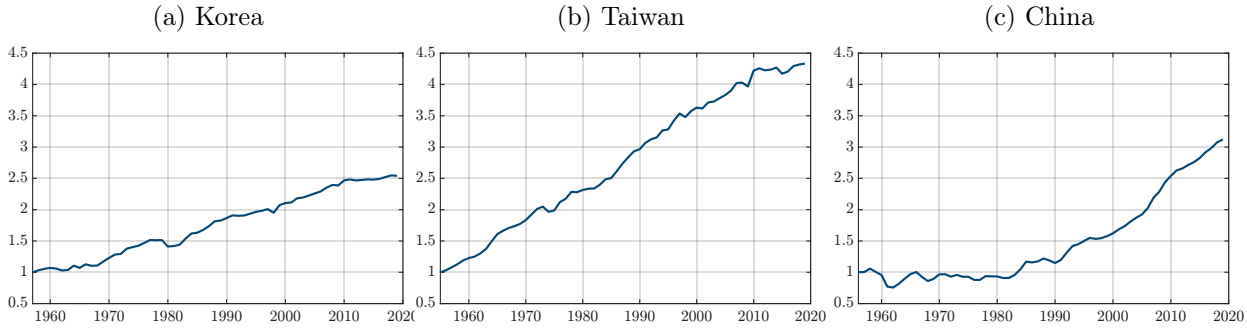

Notes: TFP at constant national prices (2017=1) for KOR, and TWN. Welfare-relevant TFP at constant national prices (2017=1) for CHN.

Figure 10: Posterior Probability, Select Countries in PWT Dataset

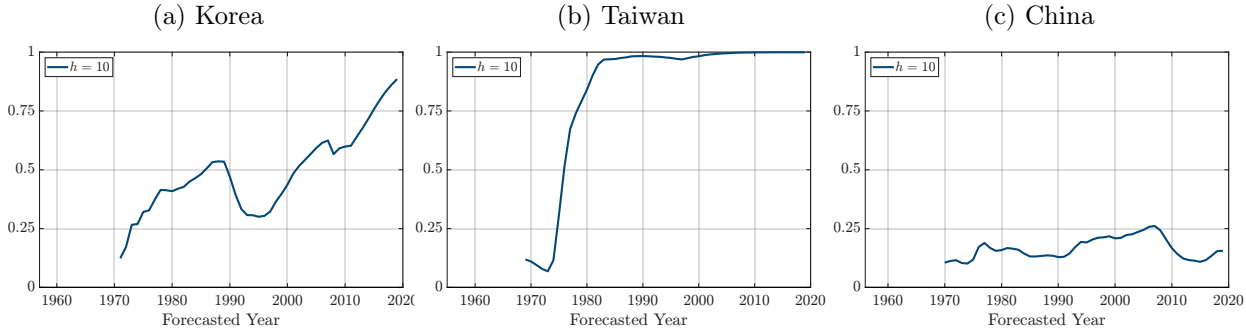

Notes: TFP at constant national prices (2017=1) for KOR, and TWN. Welfare-relevant TFP at constant national prices (2017=1) for CHN. The prior on the geometric model for each country is set to the posterior probability calculated for the U.S. in the period that the forecast is made (a value which is, for all countries considered here, close to 0.9, the prior considered for the U.S. at the start of the analysis).

Figure 11: Posterior Model Weight, 10 Year Forecast Horizon, Recursive Estimation

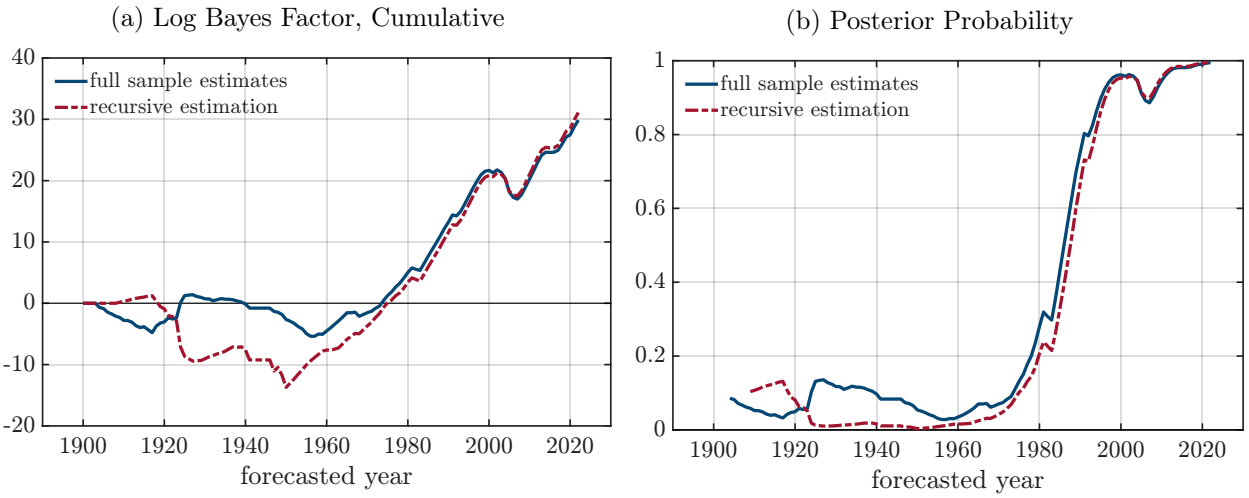

Notes: When computing the posterior probability, the initial prior of model $\mathcal{G}$ is 0.9.

We use standard Bayesian techniques and an uninformative prior. We describe the algorithm we use to estimate the parameters in Appendix B.3.

Posterior estimates. Table 2 gives the mode and standard deviation of the estimated posterior distributions. Figure B3 in the Appendix shows the posterior densities of the parameters. These moments of the posterior distributions are similar to the MLE estimates presented in the main text.

Table 2: Bayesian Parameter Estimates (US, 1880-2022)

	Model $\mathcal{A}$		Model $\mathcal{G}$	
	σ_u	σ_a	σ_ν	σ_g
Mode	0.0048	0.0681	0.0006	0.0324
Standard Deviation	(0.0037)	(0.0048)	(0.0012)	(0.0021)

Notes: This table provides the mode and standard deviation of the posterior distributions of the parameters governing the standard deviations of the $\mathcal{A}$ and $\mathcal{G}$ models estimated using Bayesian methods.

Posterior inference. For posterior inference, we use 1,000 parameter draws from the posterior distribution and conduct the same analysis as in the main text. Figure 12 plots the filtered states under those 1,000 draws, showing that the median b_t and g_t across those draws is similar to their filtered paths under the MLE estimates, while the distribution is somewhat narrow, particularly for the estimate of b_t. Next, Figure 13 shows the Bayesian model evaluation exercise for those 1,000 draws. As under the baseline framework, the

Figure 12: Bayesian Estimation Filtered States

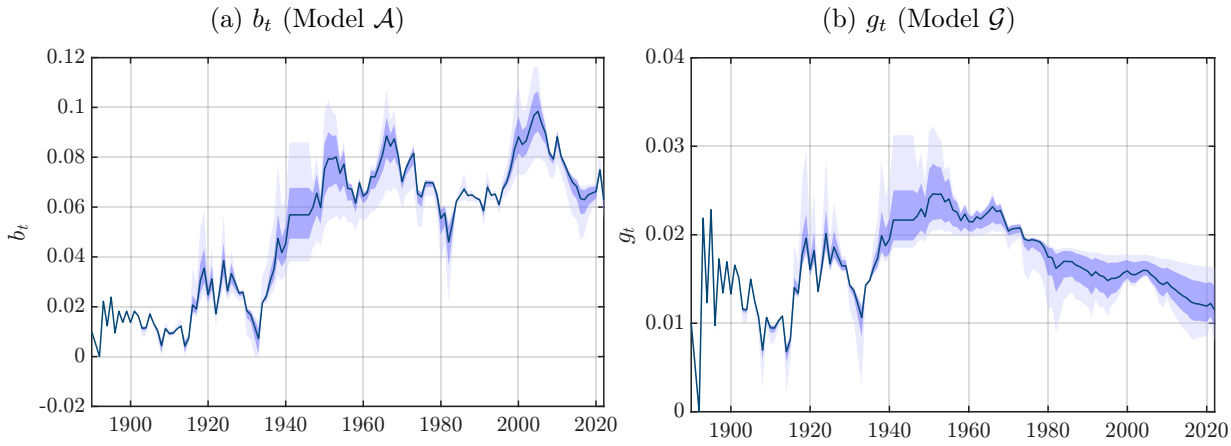

Notes: This figure shows the filtered values for b_t and g_t for 1,000 draws of the parameters from the posterior distributions. The dark blue areas show the 25th and 75th percentiles of the resulting distribution of b_t and g_t and the light blue areas show the corresponding 5th and 95th percentiles.

median posterior probability of model $\mathcal{A}$ approaches 1 by the year 2000. Examining the distribution of posterior probabilities, by 2022, the end of the sample, most draws imply the posterior model weight on model $\mathcal{A} \approx 1$. In addition, the pickup in growth around 2000 caused the posterior model weight of model $\mathcal{A}$ to decline a touch at the median across draws–as in the baseline specification–but did lead to a sharper widening of implied posterior probabilities around the year 2000.

6.3 Autoregressive Specification for Trend Growth

We next extend the baseline stochastic trend specifications by allowing the latent growth components to follow stationary autoregressive processes. This formulation nests the random-walk trend models as a limiting case and permits explicit estimation of persistence and long-run trend growth. We now specify model $\mathcal{A}$ (additive growth) as

$$A_t = A_{t-1} + b_t + \sigma_a \epsilon_t^a$$
$$b_t = (1 - \rho_b) b + \rho_b b_{t-1} + \sigma_u u_t$$

where $\rho_b \in (0, 1)$ governs persistence, b is the long-run mean growth rate, and u_t and ϵ_t^a are independent standard normal shocks, as before. Similarly, we now specify model $\mathcal{G}$

Figure 13: Model Evaluation, Bayesian Estimation

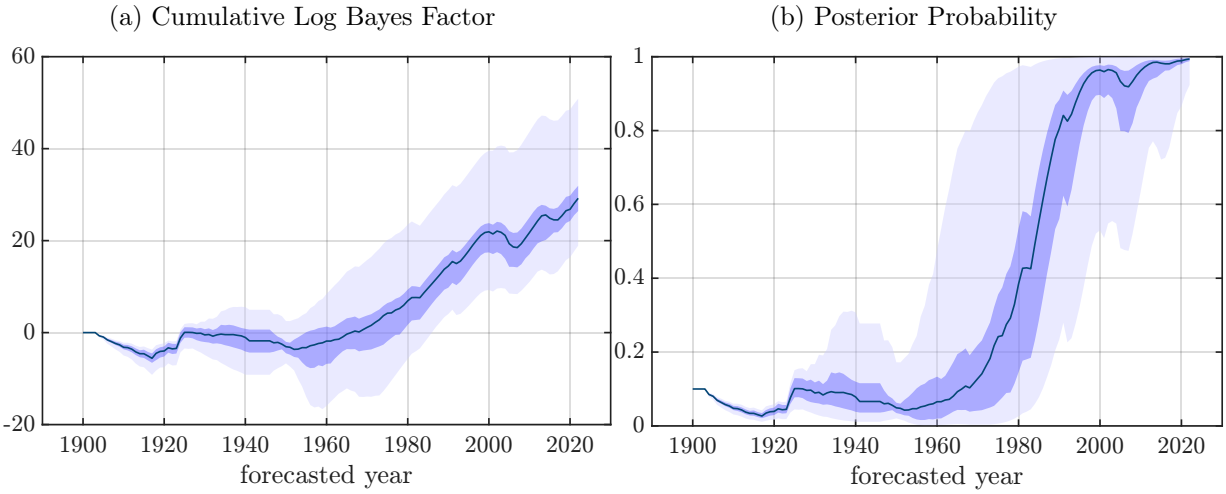

Notes: This figure shows the cumulative log Bayes factor and the posterior probability of model $\mathcal{A}$ for 1,000 draws of the parameters from the posterior distributions. The dark blue areas show the 25th and 75th percentiles and the light blue areas show the corresponding 5th and 95th percentiles. The initial prior on model $\mathcal{G}$ is 0.9.

(geometric growth) as

$$A_t = A_{t-1}\left(1 + g_t\right)\exp\left(\sigma_g \epsilon_t^g\right)$$
$$g_t = (1 - \rho_g)g + \rho_g g_{t-1} + \sigma_\nu \nu_t$$

where $\rho_g \in (0, 1)$ governs persistence, g is the long-run mean growth rate, and ν_t and ϵ_t^g are independent standard normal shocks.

We estimate each model using Bayesian Markov chain Monte Carlo methods. The posterior density combines the Kalman filter likelihood with the following priors: the persistence parameters ρ_b and ρ_g follow Beta distributions on $(0, 1)$, the innovation variances σ_u^2, σ_a^2, σ_ν^2, and σ_g^2 follow inverse-gamma priors, the long-run mean growth parameters b and g are assigned diffuse uniform priors.

Table 3 presents the mode and standard deviation of the posterior distributions. Both models $\mathcal{A}$ and $\mathcal{G}$ exhibit persistence in their mean growth parameters, more-so for model $\mathcal{A}$ (around 0.94 for model $\mathcal{A}$ and around 0.7 for model $\mathcal{G}$). Consequently, the filtered path of b_t, shown in panel (a) of Figure 14 looks very similar to our baseline random walk specification. In addition, panel (b) of the figure shows that the filtered path of g_t follows the same broad pattern as our baseline random walk specification (high in the 1940s and 1950s, before generally gradually declining through to the end of the sample), though with more noise. The conclusions regarding the Bayesian preferred model, shown in Figure 15 are similar to

Table 3: Bayesian Parameter Estimates, Autoregressive Model (US, 1880-2022)

	Model $\mathcal{A}$			
	ρ_b	σ_u	σ_a	b
Mode	0.9442	0.0064	0.0677	0.0545
Standard Deviation	(0.2031)	(0.0024)	(0.0046)	(0.0219)
	Model $\mathcal{G}$			
	ρ_g	σ_ν	σ_g	g
Mode	0.7350	0.0052	0.0313	0.0113
Standard Deviation	(0.2948)	(0.0013)	(0.0021)	(0.0046)

Notes: This table provides the mode and standard deviation of the posterior distributions of the parameters governing the standard deviations of the $\mathcal{A}$ and $\mathcal{G}$ models with the autoregressive specification estimated using Bayesian methods.

those of our baseline specification, with the posterior model weight of model $\mathcal{A} \approx 1$ towards the end of the sample.

6.4 Piecewise Regime Breaks

In this section, we explore the robustness of our results using estimations with structural breaks. We consider an estimation treating b and g as parameters to be estimated. In Appendix B, we consider a separate estimation with structural breaks in the volatility of the shocks.

6.4.1 Piecewise b and g

We estimate versions of the models in which the drift (b_t) in the additive model or the growth rate (g_t) in the geometric model follows a piecewise constant process, with a fixed number of regimes and unknown break dates. We now specify model $\mathcal{A}$ as

$$A_t = A_{t-1} + b_t + \sigma_a \varepsilon_t^a, \quad b_t = b^{(k)} \text{ for } t \in \mathcal{R}_k,$$

where ε_t^a is a standard normal shock, and model $\mathcal{G}$ as

$$A_t = A_{t-1}(1 + g_t)\exp(\sigma_g \varepsilon_t^g), \quad g_t = g^{(k)} \text{ for } t \in \mathcal{R}_k,$$

where ε_t^g is a standard normal shock. In both of these specifications, $\mathcal{R}_k$ denotes the k^{th} regime, and the parameters $\{b^{(k)}\}$ or $\{g^{(k)}\}$ are constant within regimes but allowed to differ

Figure 14: Autoregressive Model, Filtered States

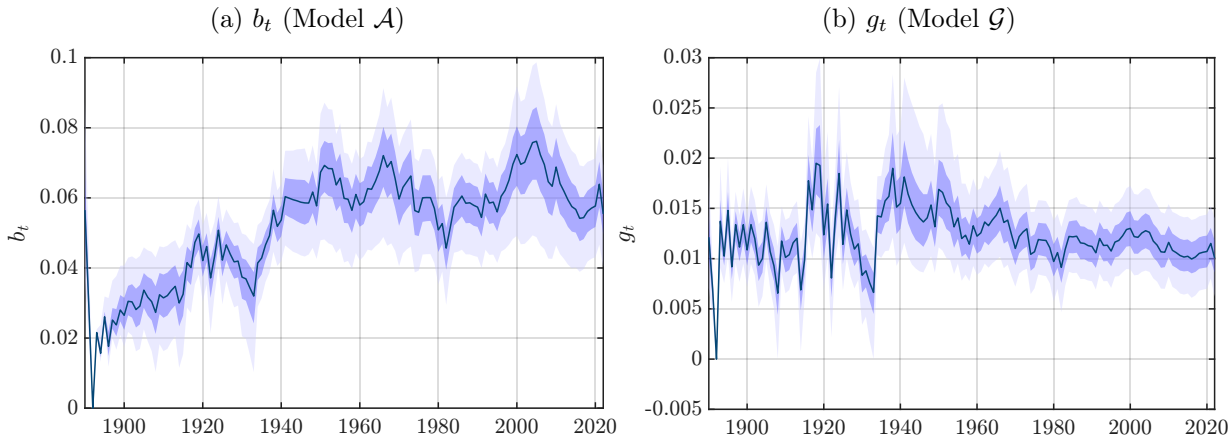

Notes: This figure shows the filtered values for b_t and g_t for 1,000 draws of the parameters from the posterior distributions. The dark blue areas show the 25th and 75th percentiles of the resulting distribution of b_t and g_t and the light blue areas show the corresponding 5th and 95th percentiles.

across them.

We consider the following specification: the number of regimes is fixed at $K_A = 3$ for Model $\mathcal{A}$ and $K_G = 3$ for Model $\mathcal{G}$ (i.e. two breakpoints per model); the innovation variances σ_a and σ_g are constant across time; regime break dates $\{t_1, \ldots, t_{K-1}\}$ are unknown and estimated; and regime lengths have a minimum length of 20 periods.

Estimation procedure. The posterior over parameters and breakpoints is explored using a Metropolis-within-Gibbs sampler, which is largely standard. It has the following steps:

1. **Initialize** the parameters $\theta = \{b^{(1)}, \ldots, b^{(K)}, \sigma_a\}$ and breakpoints.

2. **Gibbs Step:** Alternate between:

   (a) **Parameter Step:** Propose a new θ' using a random-walk Metropolis step:

   $$\theta' = \theta + \kappa \cdot \text{chol}(H^{-1}) \cdot \eta, \quad \eta \sim \mathcal{N}(0, I)$$

   where H is the negative Hessian from the initial maximum likelihood estimation in which the regimes are equally spaced over the sample.

   (b) **Breakpoint Step:** Propose new breakpoints using a local random shift mechanism that maintains the minimum separation constraint. Accept or reject the new breakpoints using a Metropolis step based on the change in log-posterior.

Figure 15: Posterior Probability of Model $\mathcal{A}$, Autoregressive Specification

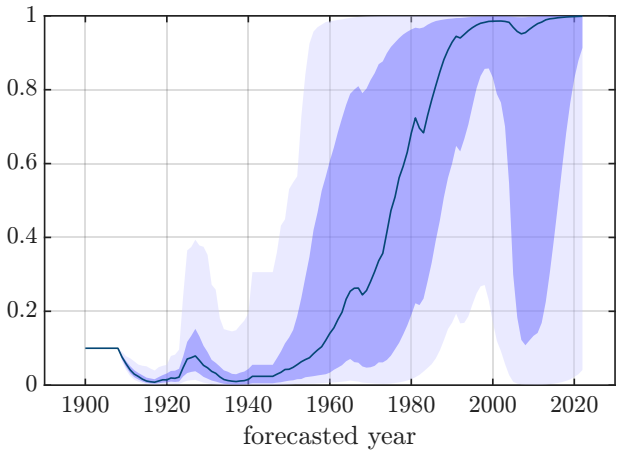

Notes: This figure shows the cumulative log Bayes factor and the posterior probability of model $\mathcal{A}$ for 1,000 draws of the parameters from the posterior distributions. The dark blue areas show the 25th and 75th percentiles of the resulting distribution of b_t and g_t and the light blue areas show the corresponding 5th and 95th percentiles. The initial prior on model $\mathcal{G}$ is 0.9.

We use uninformative priors are used on the breakpoints and independent log-normal priors are imposed on the growth parameters and the standard deviations. An adaptive tuning procedure to target an acceptance rate of around 30%. We use 1,000,000 MCMC iterations with a burn-in of 100,000 draws.

6.4.2 Estimation Results

Figure 16 plots the posterior estimates of b and g with structural breaks. The estimates for b are consistent with those of the random walk filtered values for b_t under the MLE estimates shown in the main text, with a notable break in the late 1930s, after which it remains roughly stable. The estimates for g are also roughly consistent with the estimated path of g_t under the MLE estimates shown in the main text, with a step up in the estimate value around the 1940s, and a step down around the 1960s.

Under these estimates, Figure B4 shows the posterior probability of the $\mathcal{A}$ model for 1,000 draws of the posterior distribution. Consistent with all our earlier results, we find that the posterior model weight on model $\mathcal{A} \approx 1$ by the end of the sample. In Appendix B.4, we additionally consider the case where the number of regimes in the $\mathcal{A}$ model for b is set to two (i.e. $K_A = 2$, or one break in b), and compare that to the case where there are three regimes for g for Model $\mathcal{G}$ (two breaks in g). This choice allow for more flexibility for the geometric model compared to the additive model. As shown in the Appendix, under this specification, we find that the posterior model weight on model $\mathcal{A} \approx 1$ by the end of the sample.

Figure 16: Piecewise b and g Estimates

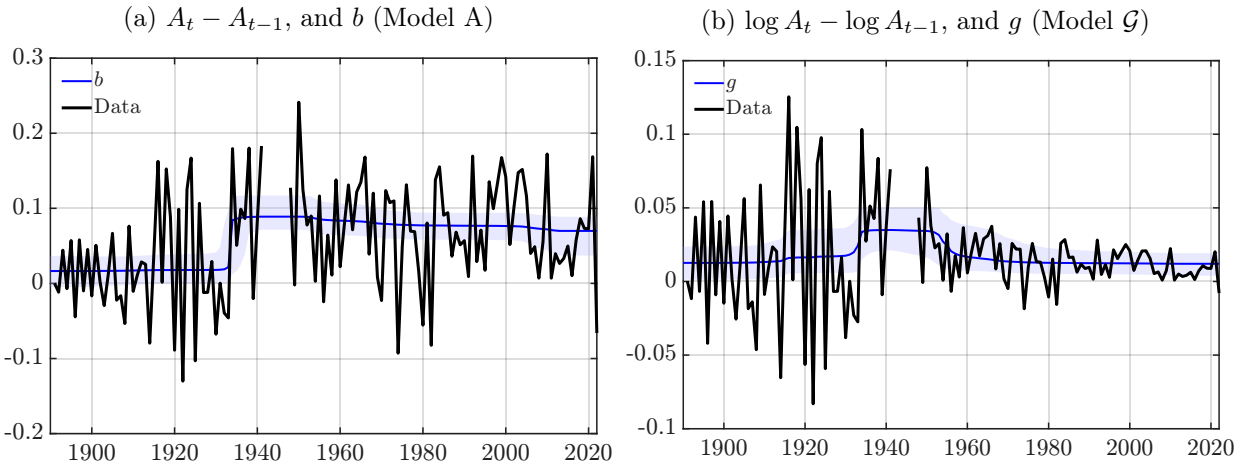

Notes: Both models allow for two breaks in b and g. The data is plotted in black. The median of the estimate of b and g over time is plotted as the blue line. The shaded area shows the 5% and 95% posterior bands.

7 Conclusion

U.S. TFP has been growing linearly over the past 90 years, and the additive model beats the exponential model for most countries – developed or developing – where long data series of TFP are available.

The additive model, unlike the exponential one, provides useful long-term forecasts of TFP. Probably not coincidentally, professional forecasters use models whose forecasts are consistent with those of the additive model.

The additive growth model explains the observed TFP slowdown as a consequence of model misspecification. We should not have expected growth *rates* to be constant in the first place. Additive TFP growth predicts *increasing* increments of labor productivity, wages, and GDP per capita thanks to capital accumulation.

Philippon (2022) discusses further historical evidence as well as implications for economic theory. The TFP frontier appears to grow linearly within broad historical periods: 1650 to 1830, 1830 to 1930, and 1930 until today. The early break point around 1650 is consistent with Bouscasse et al. (2021).

Additive growth has implications for macroeconomic, industry and firms dynamics. Wealth effects of TFP news are weaker and long run consumption risk is lower. Additive growth also matters for the optimal mitigation of climate change, since real discount rates are low and future generations will not be much richer than the current one.

At the firm level, Lenzu et al. (2023) finds that productivity grows linearly with the age of a firm. More importantly, the study of industries and firms can shed light on *why* growth

Figure 17: Posterior Weight of Model $\mathcal{A}$ under Regime Breaks Specification

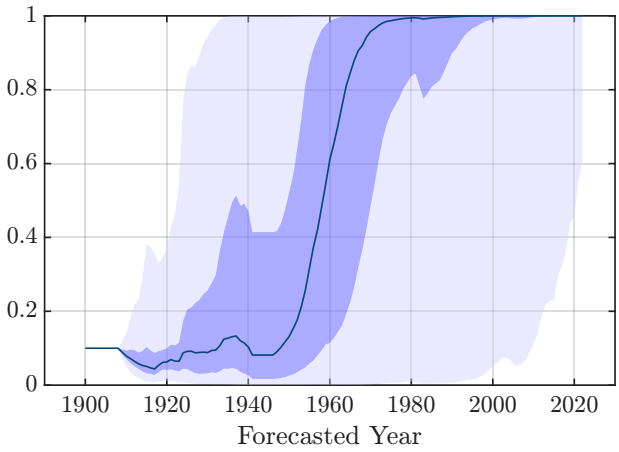

Notes: Both the $\mathcal{A}$ and $\mathcal{G}$ models allow for two breaks in the trend growth parameters b and g. The median is plotted as the blue line. The light blue shaded area shows the 5% and 95% posterior bands in light blue, the darker blue shaded area shows the 25% and 75% posterior bands.

is additive and *where* knowledge transfers take place.

At a more theoretical level, additive growth suggests that new ideas add to our stock of knowledge but do not multiply it. This observation does not rule out positive spillover from R&D across firms, and especially across vintages of firms. The additive model also does not solve the research productivity puzzle, which is not about the stochastic process for TFP but rather about the specification of the production function for ideas.

References

Akcigit, U. and S. T. Ates (2023): "What Happened to US Business Dynamism?" *Journal of Political Economy*, 131, 2059–2124.

Bergeaud, A., G. Cette, and R. Lecat (2016): "Productivity Trends in Advanced Countries between 1890 and 2012," *Review of Income and Wealth*, 62, 420–444.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2020): "Are Ideas Getting Harder to Find?" *American Economic Review*, 110, 1104–44.

Bouscasse, P., E. Nakamura, and J. Steinsson (2021): "When Did Growth Begin? New Estimates of Productivity Growth in England from 1250 to 1870," Working Paper.

Comin, D., W. Easterly, and E. Gong (2010): "Was the Wealth of Nations Determined in 1000 bc?" *American Economic Journal: Macroeconomics*, 2, 65–97.

David, P. A. (1990): "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *The American Economic Review*, 80, 355–361.

Del Negro, M., R. B. Hasegawa, and F. Schorfheide (2016): "Dynamic prediction pools: An investigation of financial frictions and forecasting performance," *Journal of Econometrics*, 192, 391–405, innovations in Multiple Time Series Analysis.

Edge, R., T. Laubach, and J. C. Williams (2007): "Learning and shifts in long term productivity growth," *Journal of Monetary Economics*, 54, 2421–2438.

Fernald, J. G. (2012): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," FRBSF Working Paper (updated March 2014).

Field, A. J. (2003): "The Most Technologically Progressive Decade of the Century," *American Economic Review*, 93, 1399–1413.

Gordon, R. J. (2016): *The Rise and Fall of American Growth*, Princeton University Press.

Guzey, A., E. Rischel, Applied Divinity Studies, Anonymous, and Anonymous (2021): "Issues with Bloom et al's "Are Ideas Getting Harder to Find?" and why total factor productivity should never be used as a measure of innovation," https://guzey.com/economics/bloom/.

Jones, B. F. (2009): "The Burden of Knowledge and the Death of the Renaissance Man: Is Innovation Getting Harder?" *The Review of Economic Studies*, 76, 283–317.

Jones, C. I. (1995): "RD-Based Models of Economic Growth," *Journal of Political Economy*, 103, 759–784.

——— (2002): "Sources of U.S. Economic Growth in a World of Ideas," *American Economic Review*, 92, 220–239.

Lenzu, S., T. Philippon, and J. Tielens (2023): "Additive Firm Dynamics," NYU WP.

Philippon, T. (2022): "Additive Growth," Working Paper.

Solow, R. M. (1956): "A Contribution to the Theory of Economic Growth," *The Quarterly Journal of Economics*, 70, 65–94.

——— (1957): "Technical Change and the Aggregate Production Function," *The Review of Economics and Statistics*, 39, 312–320.

Appendix

For Online Publication

A US Data

A.1 Three Measures of Post-War US TFP

Figure A1 compares the TFP series from BCL and Fernald, with and without adjustment for education.

Figure A1: US TFP Levels

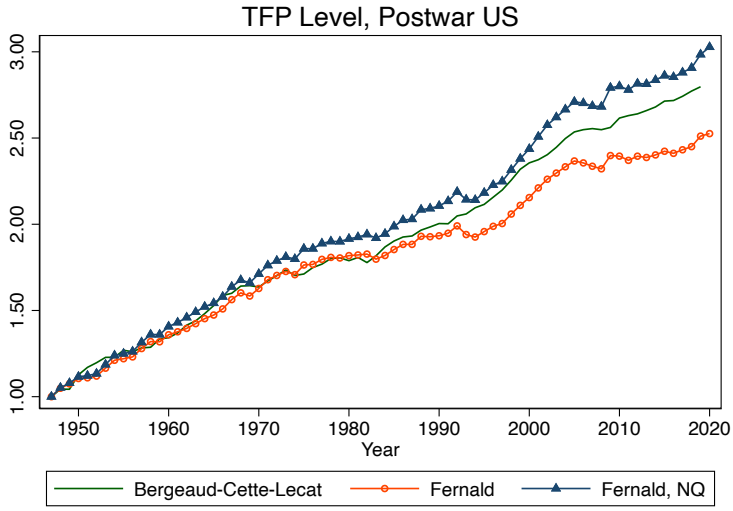

Notes: TFP levels, A^{bcl}, A^f_t, and A^n_t. Data from Fernald (2012) and Bergeaud et al. (2016).

A.2 Survey of Professional Forecasters

The data from the Survey of Professional Forecasters is summarized in the following table.

Variable Name	Description	Date Availability
PROD10	Quarterly forecasts of annual-average productivity growth rates, where rates are annualized and productivity is measured by output per hour, from the survey date to 10 years later (40 quarters total). Obtained from Survey of Professional Forecasters	1992:Q1 to 2022:Q3

A.3 Council of Economic Advisors

This section summarizes the projection data obtained from the Council of Economic Advisors (CEA).

The CEA projections for output per hour growth are taken from the annual Economic Report of the President. These forecasts are consistently available since 1983, with the exception of 1995 and 2013. Forecasts are also available in the 1970 and 1971 Reports, but with indefinite time horizons. While the values for 1970 and 1971 must be inferred from the text of the Report itself, all other values are gathered from tables in the report. The numbers of these tables in each report are listed in the supplemental data. These tables are titled the following:

- 1983: "Projections of Economic Goals,"
- 1984 to 1991: "Administrative Economic Assumptions,"
- 1992 to 2005: "Accounting for Growth in Real GDP,"
- 2006 to 2009: "Supply-Side Components of Real GDP Growth,"
- 2010 to 2011: "Components of Potential Real GDP Growth,"
- 2012: "Components of Actual and Potential Real GDP Growth,"
- 2014 to 2022: "Supply-Side Components of Actual and Potential Real GDP Growth."

Time horizons of the forecasts fluctuate over the course of the data. Between 1983 and 1994, these forecasts have a six-year time horizon, including the year of the report. From 1996 to 2002, the forecasts' time horizons range from seven to eleven years out, before returning to a six-year forecast time horizon until 2009. Following 2009, the forecasts consistently have an eleven year time horizon, with the exception of the 2013 Report with no forecasts.

The granularity of the forecasts also varies over the reports. The 1970 and 1971 reports forecast 2.8 and 3 percent annual growth of output per hour in perpetuity. Meanwhile, between 1983 and 1991, each projected year receives its own distinct forecast value. However, after 1991, all Reports give only a single per annum forecast value, always over a period that includes at least a year–if not more–of realized values.

B Additional Results using U.S. Data

B.1 Point Forecast Evaluation

This section complements the density-based model comparison in the main text by evaluating the point forecast performance of the $\mathcal{A}$ and $\mathcal{G}$ models. The goal is to isolate differences in conditional mean dynamics.

As discussed in Section 3, the predictive density comparison jointly reflects differences in both the mean and variance implied by the two models. Here, we conduct a point forecast evaluation based on mean squared error (MSE). For each model, we construct h-step-ahead point forecasts using the conditional mean implied by the estimated state-space representation. Under the additive model, the h-step-ahead forecast for the level of the series is given by $A_{t+h} = A_t + h\, b_t$, where b_t denotes the filtered drift. Under the geometric model, the corresponding level forecast is $A_{t+h} = A_t \exp(h\, g_t)$, where g_t is the filtered growth rate. In both cases, forecasts are evaluated against realized levels, and forecast errors are computed in levels.

Using the point forecasts implied by each model's conditional mean, we compute mean squared forecast errors for an 10-year-ahead horizon. The $\mathcal{A}$ model yields an MSE of 0.068, compared with 0.158 for the $\mathcal{G}$ model. In root mean squared error terms, the $\mathcal{A}$ model has a RMSE of 0.26, while the $\mathcal{G}$ model has a RMSE of 0.40. The resulting MSE ratio of about 0.43 indicates that the additive model reduces mean squared forecast error by more than half relative to the geometric alternative.

B.2 Other Forecast Horizons

The baseline forecast horizon used to evaluate the $\mathcal{A}$ and $\mathcal{G}$ models in the Bayesian model evaluation is 10 years. This section instead presents the Bayesian model evaluation under 5 and 15 year forecast horizons, alongside the baseline 10 year forecast horizon for comparison. Figure B2 shows the cumulative log Bayes factor and the posterior probability. The results are consistent with the 10-year forecast evaluation: by the end of the sample, the $\mathcal{A}$ model is preferred with a posterior weight ≈ 1. The model evaluation also fluctuates less around the prior at shorter horizons before the 1970s.

B.3 Bayesian Estimation of the Parameters

The algorithm we use is standard and is described next.

Figure B2: Posterior Model Probability, $h = 5, 10, 15$ Year Forecast Horizons

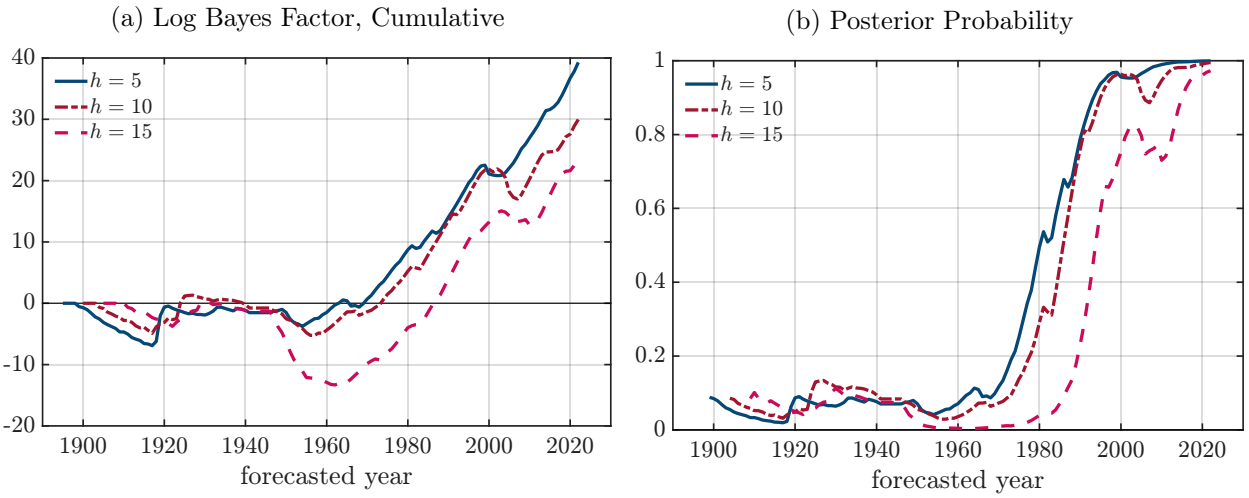

Notes: When computing the posterior probability, the initial prior of model $\mathcal{G}$ is 0.9.

Metropolis-Hastings algorithm. Given current parameter vector θ, a new proposal θ' is drawn using a random walk based on a Student-t proposal with 12 degrees of freedom:

$$\theta' = \theta + \kappa \cdot \text{chol}\left(H^{-1}\right) \cdot \varepsilon, \quad \varepsilon \sim t_{12}\left(0, I\right),$$

where H is the Hessian from the log-likelihood function evaluated at the maximum likelihood estimate, and κ is a tuning constant for step size. The acceptance probability is computed as:

$$\alpha = \min\left(1, \exp\left[\log p\left(y \mid \theta'\right) + \log \pi\left(\theta'\right) - \log p\left(y \mid \theta\right) - \log \pi\left(\theta\right)\right]\right),$$

where $p\left(y \mid \theta\right)$ is the likelihood (evaluated using the Kalman filter), and $\pi\left(\theta\right)$ is the prior. During the burn-in period, the scaling factor κ is updated every 1000 iterations to target an acceptance rate of 30%.

Figure B3 plots the posterior densities of the estimated parameters, showing they are single peaked and their modal values are close to those of the baseline MLE estimation.

B.4 Estimations with Structural Breaks

B.4.1 Regime Breaks in Trend Growth

This section contains additional results for the estimation of the model with regime breaks in the growth parameters b and g. Figure B4 shows the posterior probability of the $\mathcal{A}$ model under the piecewise break specification, allowing for two breaks in b in the $\mathcal{A}$ model and three breaks in g in the $\mathcal{G}$ model. Consistent with the results shown in the main text, under

Figure B3: Posterior Density of Estimated Parameters

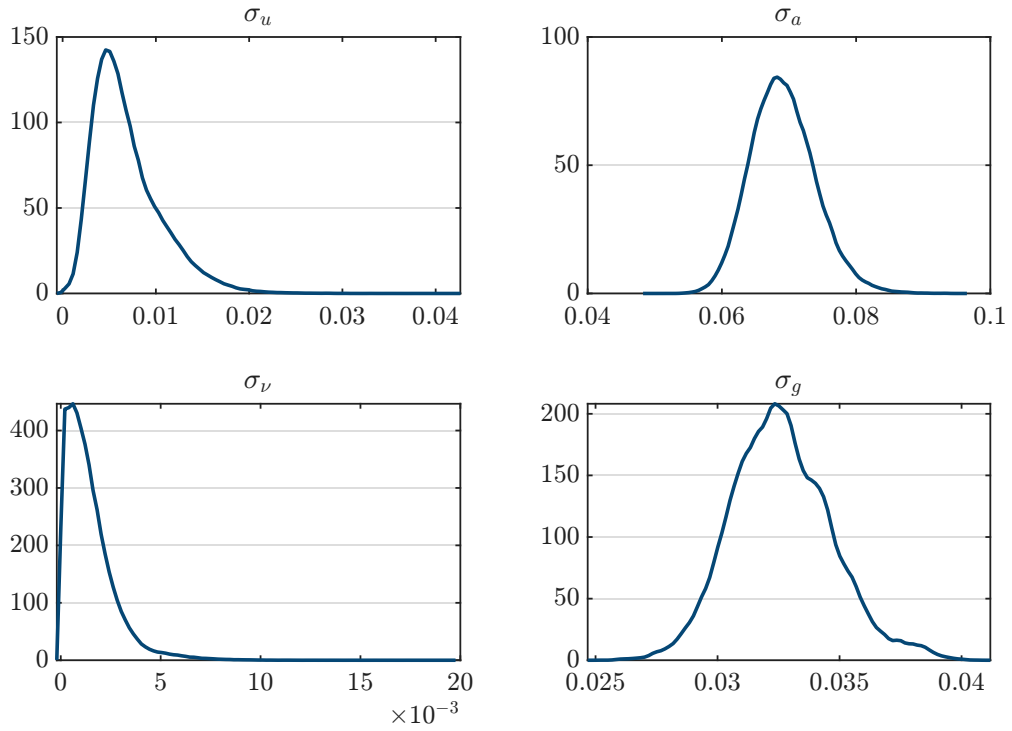

Notes: This figure shows the posterior distributions of the parameters governing the standard deviations of the $\mathcal{A}$ and $\mathcal{G}$ models estimated using Bayesian methods..

Figure B4: Posterior Probability Weight of $\mathcal{A}$ Model, Piecewise Model, $K_{\mathcal{A}} = 2$, $K_{\mathcal{G}} = 3$.

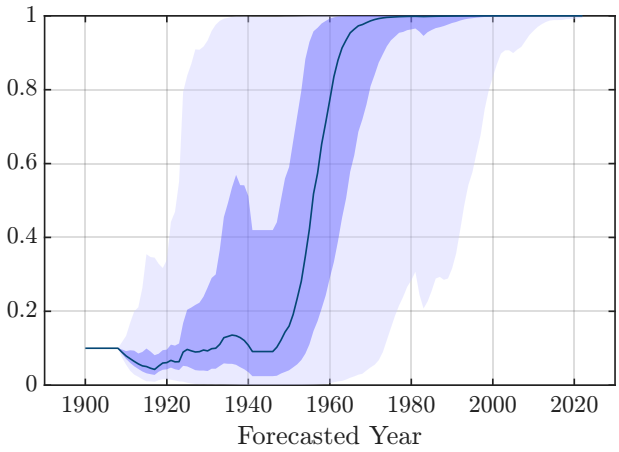

Notes: The median is plotted as the blue line. The light blue shaded area shows the 5% and 95% posterior bands in light blue, the darker blue shaded area shows the 25% and 75% posterior bands.

these specifications we find that the posterior model weight of model $\mathcal{A} \approx 1$ towards the end of the sample.

B.4.2 Posterior Estimates of Baseline Model with Regime Breaks

Figures B5 and B6 show the posterior distributions of the parameters under the piecewise specifications for b and g in which we allow for three regimes in either model. Panel (a) shows the estimate values of b and g across the three regimes, while Panel (b) shows the posterior distributions of the breakpoint dates for the regime breaks. For model $\mathcal{A}$, there is a clear increase in the estimate of b between the first and second regimes around the 1930s, with less evidence of a structural break between the second and third regimes (reflected in a wide posterior of the break date between b^2 and b^3). Model $\mathcal{G}$ shows more evidence of a step-down in geometric growth around the 1950s to 1960s.

B.4.3 Time-varying Variances

We next estimate the two models with time-varying stochastic volatility.

Additive model:

$$A_t = A_{t-1} + b_t + \sigma_{a,t}\,\varepsilon_t^a,$$
$$b_t = b_{t-1} + \sigma_u\,u_t,$$

Figure B5: Posterior Densities, Model $\mathcal{A}$, Model with Regime Changes

(a) Estimated Parameters

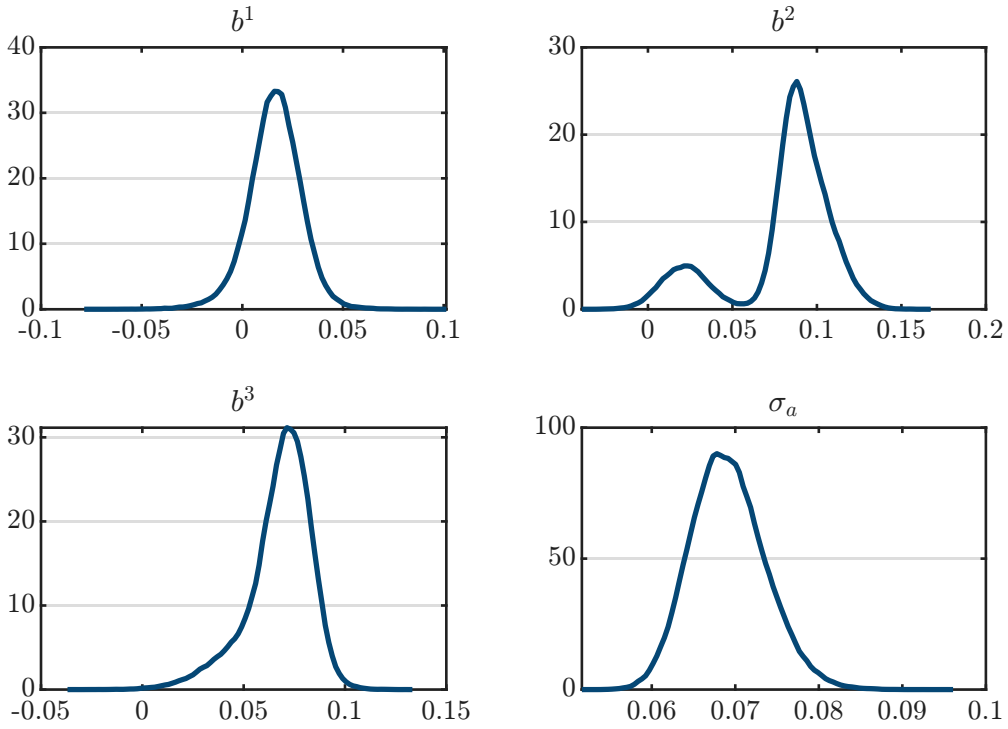

(b) Estimated Break Dates

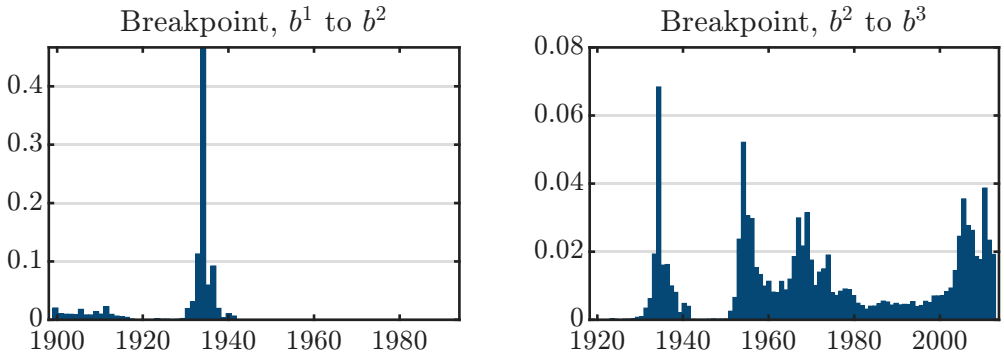

Notes: This figure shows the posterior distributions for the parameters of the $\mathcal{A}$ model with structural breaks in b.

Figure B6: Posterior Densities, Model $\mathcal{G}$, Model with Regime Changes

(a) Estimated Parameters

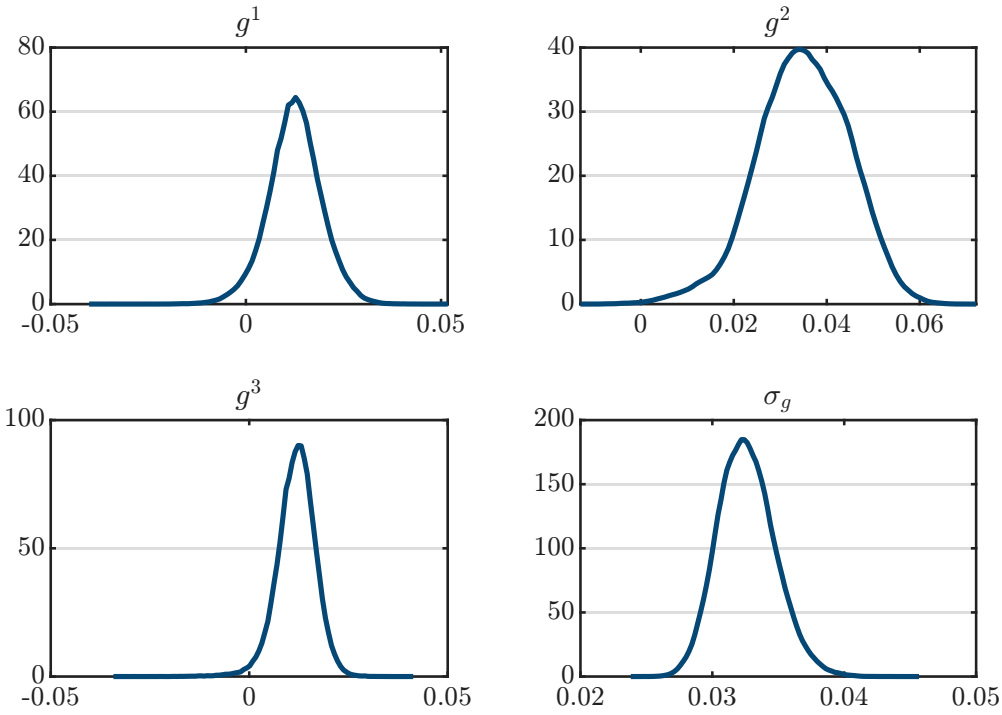

(b) Estimated Break Dates

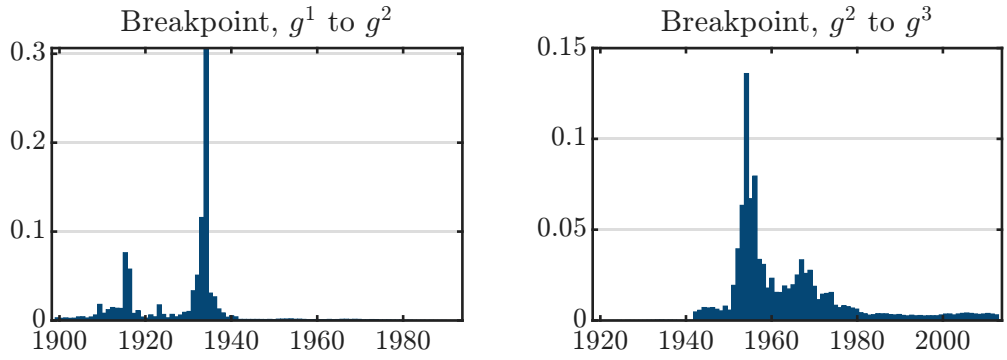

Notes: This figure shows the posterior distributions for the parameters of the $\mathcal{G}$ model with structural breaks in g.

Figure B7: Posterior Probability of Model $\mathcal{A}$, Regime Breaks in Volatility

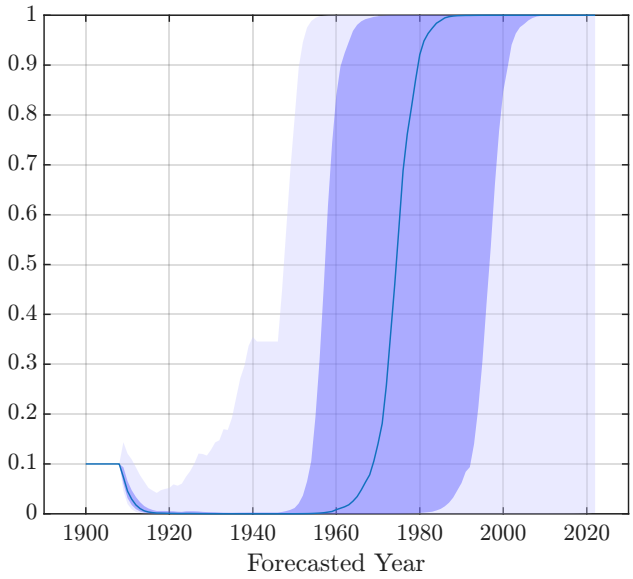

Notes: Based on 1,000 posterior draws. Both models allow for four regimes, or three breaks, in σ_a or σ_g. The median is plotted as the blue line. The light blue shaded area shows the 5% and 95% posterior bands in light blue, the darker blue shaded area shows the 25% and 75% posterior bands.

where u_t and ε_t^a are independent standard normal shocks. The innovation variance in the TFP equation, $\sigma_{a,t}^2$, is allowed to change over time via a piecewise-constant structure. We assume K regimes, with regime-specific values for $\sigma_{a,t}$:

$$\sigma_{a,t} = \sigma_a^{(k)} \quad \text{for } t \in \text{Regime } k.$$

Geometric model:

$$A_t = A_{t-1}(1 + g_t)\exp(\sigma_{g,t}\,\varepsilon_t^g),$$
$$g_t = g_{t-1} + \sigma_\nu\,\nu_t,$$

where ν_t and ε_t^g are standard normal shocks. As in the additive model, the innovation standard deviation $\sigma_{g,t}$ is allowed to vary across regimes:

$$\sigma_{g,t} = \sigma_g^{(k)} \quad \text{for } t \in \text{Regime } k.$$

We consider the case of four regimes in the volatilities, and use a similar estimation procedure as above for the case of breaks in b and g.

Figure B7 shows the posterior probability weight of model $\mathcal{A}$ under this specification, showing as above that the weight ≈ 1 towards the end of the sample.

B.5 Labor Productivity

Let us now study the accumulation of capital. Define the capital labor ratio as

$$k_t \equiv K_t / L_t,$$

where, in the BCL data, K_t is the real capital stock and L_t measures hours worked. The first order condition for capital demand in the neoclassical growth model equates the marginal product of capital (MPK) to the user cost (defined as χ). BCL do not consider changes in the user cost and the first order condition is simply

$$k_t^{1-\alpha} = \frac{\alpha}{\chi} A_t. \tag{9}$$

Equation (9) says that the normalized inverse MPK (IMPK) is proportional to A.[6] Model $\mathcal{G}$ therefore predicts that $k_t^{1-\alpha}$ grows exponentially, while model $\mathcal{A}$ says that it grows linearly. Once we have a forecast for the capital labor ratio we can use our forecast for TFP to create a forecast for labor productivity λ_t, defined as output per hour:

$$\lambda_t \equiv \frac{Y_t}{L_t} = A_t k_t^{\alpha}. \tag{10}$$

model $\mathcal{A}$ offers a forecast for labor productivity as

$$\hat{\lambda}_t = \left(\hat{a} + \hat{b}t\right)\left(\hat{a}_{impk} + \hat{b}_{impk}t\right)^{\frac{\alpha}{1-\alpha}}$$

C International Data and Additional Results

Figure C9 plots the TFP data for all countries in the BCL dataset. Figure C10 shows the cumulative log Bayes factor for each of the countries in the BCL dataset. Finally, Table C1 gives the ratio of a simple RMSE forecast evaluation of the two models at the 10-year horizon, given each model's country estimates. With the exception of Ireland, these results show the $\mathcal{A}$ model provides superior forecast performance across all countries.

[6] Users of model $\mathcal{G}$ typically interpret equation (9) as saying that capital grows exponentially, just like A, as a rate $(1+g)^{1/(1-\alpha)}$. Equivalently, if the model is written with Harrod-neutral technological progress, $Y_t = K_t^{\alpha}\left(\mathcal{M}_t H_t\right)^{1-\alpha}$ then capital is proportional to $\mathcal{M}_t$.

Figure B8: Out-of-Sample Labor Productivity Forecasts

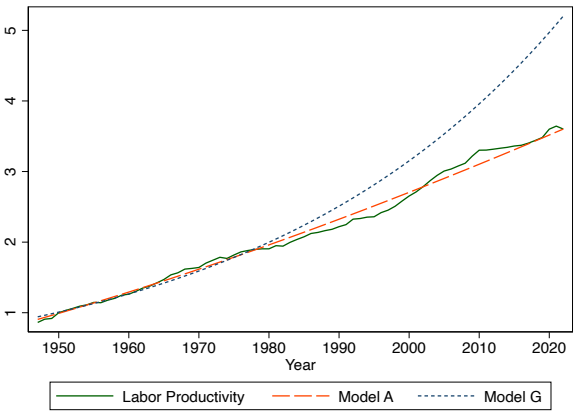

Notes: Models are estimated over 1947-1983. The forecast 1984-2019 is out-of-sample. Labor productivity is real GDP per hour. Data source: Bergeaud et al. (2016).

Figure C9: TFP Levels

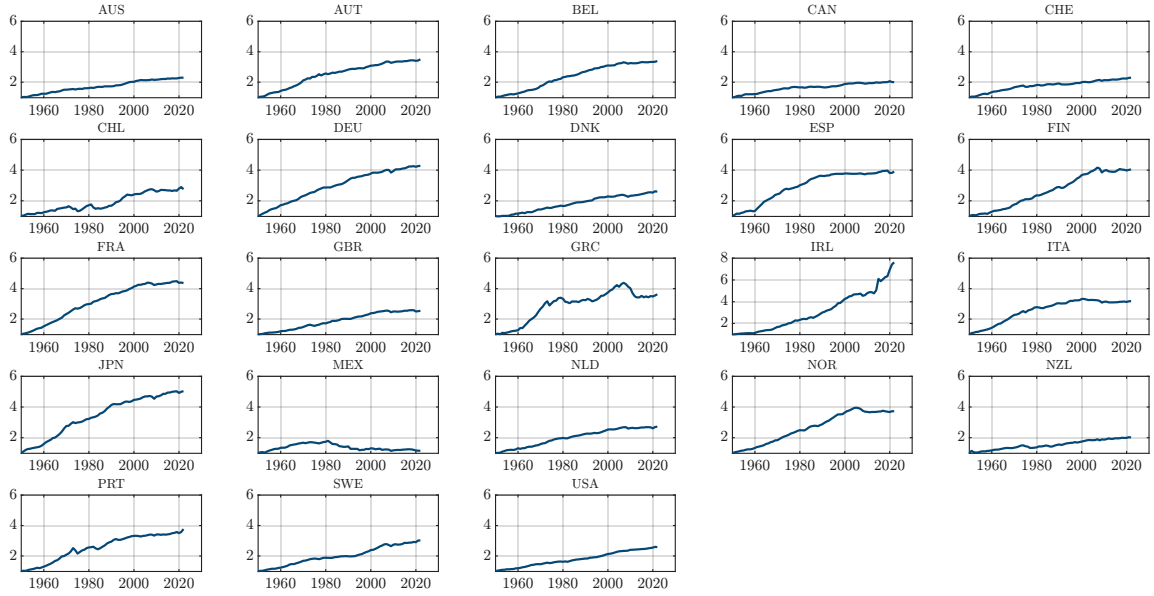

Notes: TFP levels for each country (1950=1). Data from Bergeaud et al. (2016).

Figure C10: Cumulative Log Bayes Factor

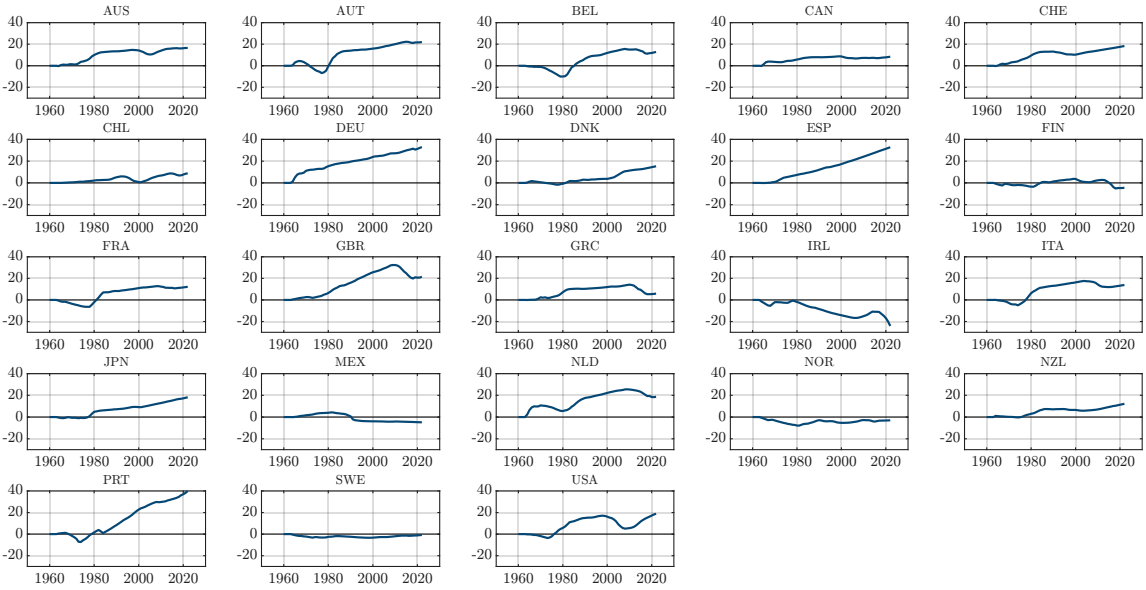

Notes: The figure shows the cumulative log Bayes factor for each country given the full-sample MLE estimates for each country.

Figure C11: Posterior Weight on Model $\mathcal{A}$, BCL Sample

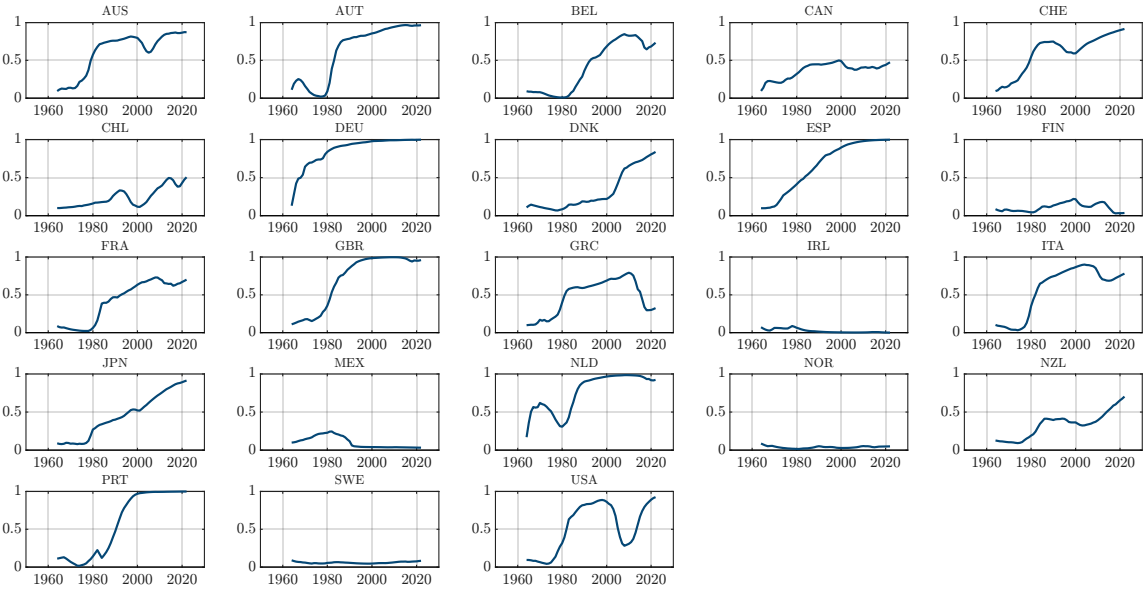

Notes: The figure shows the posterior probability weight for each country given the full-sample MLE estimates for each country.

Table C1: Ratio of RMSE of $\mathcal{A}$ Model to $\mathcal{G}$ Model, BCL Dataset

Country	RMSE($\mathcal{A}$) / RMSE($\mathcal{G}$)	Country	RMSE($\mathcal{A}$) / RMSE($\mathcal{G}$)
AUS	0.583	GRC	1.592
AUT	0.460	IRL	0.577
BEL	0.580	ITA	0.462
CAN	0.632	JPN	0.808
CHE	0.595	MEX	0.541
CHL	0.560	NLD	0.733
DEU	0.366	NOR	0.523
DNK	0.518	NZL	0.258
ESP	0.412	PRT	0.734
FIN	0.641	SWE	0.524
FRA	0.557	USA	0.572
GBR	0.446		

Notes: This table provides the ratio of the RMSE of the $\mathcal{A}$ model to the RMSE of the $\mathcal{G}$ model for the models estimated over the full sample.