NBER WORKING PAPER SERIES

ADDITIVE GROWTH

Thomas Philippon

Working Paper 29950
http://www.nber.org/papers/w29950

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2022, Revised May 2022

I am grateful to Chad Jones, Virgiliu Midrigan, Greg Mankiw, Olivier Blanchard, Xavier Gabaix, Xavier Jaravel, Antonin Bergeaud, Remy Lecat, David Weil, Gilbert Cette, David Romer, Callum Jones, German Gutierrez, Ben Jones, Alexey Guzey, and Bill Easterly for their comments. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Additive Growth
Thomas Philippon
NBER Working Paper No. 29950
April 2022, Revised May 2022
JEL No. E22,N1,O11,O3,O4

ABSTRACT

Growth theory is based on the assumption of exponential total factor productivity (TFP) growth. Across countries and time periods I find that TFP growth is actually linear. The additive growth model, unlike the exponential one, provides useful long-term forecasts for TFP. For the distant past the model suggests piecewise linear evolutions where the size of TFP increments changes in the late 1600's, the early 1800's, and around 1930. For the distant future the model predicts ever increasing increments in standards of living but with falling real interest rates and growth rates that converge to zero. The model suggests stable TFP growth in the US, but a TFP slowdown in the Euro area since the late 1990s.

Thomas Philippon
New York University
Stern School of Business
44 West 4th Street, Suite 9-190
New York, NY 10012-1126
and NBER
tphilipp@stern.nyu.edu

Following Solow (1956) textbook models of economic growth assume that TFP growth is exponential: $dA_t = gA_t dt$, where A is TFP and g is constant or at least highly persistent. I examine data across many countries and time periods and I find that, in nearly all cases, productivity growth is in fact linear: $dA_t = bdt$ where b is constant, at least within broad historical periods.

I start my investigation with post war US data from Fernald (2012) – "Fernald" – and Bergeaud et al. (2016)– "BCL". I find that TFP growth is linear in the US during the post war period. Using BCL's data US TFP 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 0.0245, i.e., 2.45% of its initial value, not 2.45% per year. Using Fernald's data for the *private* sector, the same statement holds with an annual increments of 0.0276. The size of the increments does not appear to change over 80 years. There is no TFP slowdown, or, to put it differently, the perceived TFP slowdown comes from a misspecified benchmark. Initial trend growth is around 2.5%. After 40 year, TFP doubles, and since increments are constant, the trend growth rate is half of what it used to be. After 60 years later, it is only one percent, and so on.

Linear TFP growth does not imply linear GDP growth. Capital accumulation creates a convex time path for labor productivity and GDP per capita under Hicks-neutral linear TFP growth. Empirically, the linear TFP model predicts the correct non-linear evolution of labor productivity while the exponential model over-predicts the level of future labor productivity.

In Section 2 I estimate models with time varying trend growth in the long BCL sample of 23 countries over 129 years. I find that the additive model predicts TFP dynamics better than the exponential model for each of the 23 countries. The 10-year forecast errors of the exponential model are 30% to 60% higher than those of the additive model. I also consider a sample of OECD countries that are not in the BCL sample (e.g., Korea) and I show that their TFP growth is linear. TFP growth paths in Thailand and Taiwan, two prime example of "miracle growth" in Asia, are also linear. The exponential growth model fails because it predicts periods of sustained and convex productivity growth that simply do not exist in the data.

The main contribution of this paper is empirical but Section 3 discusses some theoretical implications. Economic theory does not predict the shape of TFP growth. Models of endogenous growth such as Romer (1986b), Lucas (1988) and Aghion and Howitt (1992) *assume* a knowledge production function that implies exponential growth. For instance, models with vertical differentiation *assume* that the quality ladder is exponential. My results suggest instead that the ladder is linear and that inter-temporal spillovers are smaller than previously thought: an improvement in TFP raises economic efficiency but does not imply that future discoveries become exponentially easier. The long term properties of the additive growth model are straightforward. For a given increment b the linear model converges to a balanced growth path with a constant capital/output ratio. The capital labor ratio, labor productivity, and GDP per capita grow indefinitely, and with increasing increments. The model therefore does not predict stagnation: incomes are increasing ever faster even as growth *rates* tend to zero.

Finally, Section 4 steps away from a purely econometric approach to provide a broader historical interpretation of the data. A symptom of the failure of the exponential model is that the estimated trend growth rates are unstable. By contrast, the additive TFP model displays very few breaks and, in most cases, these breaks have a plausible economic interpretation in terms of General Purpose Technologies (GPTs). For example, the process of US TFP increments has only one break over the past 130 years, around 1930, following the large-scale implementation of the electricity revolution (Gordon, 2016). I investigate growth before 1890 using UK GDP per capita and I find two breaks between 1600 and 1914. The first is somewhere between 1650 and 1700, when growth becomes positive. The second is around 1830. These breaks are consistent with historical research on the first and second industrial revolutions (Mokyr and Voth, 2010). These rare breaks represent the main source of convexity in the historical series and TFP growth appears to be linear between the breaks.

Literature This papers sheds light on existing puzzles in the growth literature. The perceived TFP slowdown is the result of a misspecified model, since growth was never actually exponential. The additive model, unlike the exponential one, provides useful benchmarks and forecasts across a wide variety of countries and time periods. Changes in the size of TFP increments seem to happen mostly around the discovery of new GPTs. The idea that growth increments increase during industrial revolutions speaks to the complementarity of new inventions with existing technologies emphasized by Comin et al. (2010). Jones (2009) and Bloom et al. (2020) 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.

1 Evidence from US Growth

My main sources of data – Fernald (2012) and Bergeaud et al. (2016) – assume a Cobb-Douglas production function, so I will do the same in most of the paper. Aggregate value value added (GDP, Y_t) is given by

$$Y_t = A_t K_t^{\alpha} L_t^{1-\alpha}, \tag{1}$$

where A_t is TFP, K_t is the flow of capital services and L_t is the flow of labor services. Section 3 discusses more general production functions $F(K, L, A)$ and compares Hicks and Harrod neutrality in the context of linear growth. My goal is to understand the long-term dynamics of TFP. Since at least Solow (1956) economists have assumed that A follows a geometric process, which I call model G[1]:

$$\mathbb{E}[A_{t+\tau} \mid A_t] = A_t (1+g)^{\tau}. \tag{2}$$

I will show that growth is additive and that the TFP process is better described by model D (as in "difference"):

$$\mathbb{E}[A_{t+\tau} \mid A_t] = A_t + b\tau, \tag{3}$$

where b is a parameter that measures the size of increments. I start my investigation with post-war US data. The empirical justification is that this is the most widely used and reliable data. The theoretical justification is that one might expect different TFP dynamics between countries at the frontier and countries catching up to the frontier. The main advantage of post-war US data, then, is that one can reasonably argue that the US was at the technological frontier during the entire period. For the same reason I will focus on the UK when studying productivity before 1900.

1.1 Main Data Sources

My primary sources for TFP are Fernald (2012) (Fernald) and Bergeaud et al. (2016) (BCL). Let A_t^{BCL} and A_t^F denote the BCL and Fernald measures of TFP. There are several differences between these two datasets. BCL covers 23 countries from 1890 to 2019 and their data allow the analysis of a long sample as well as international

[1]Jensen's inequality terms do not play a significant role in the empirical analysis of this section.

comparisons in Section 2. Fernald's series cover only the US business sector, while BCL include households and the government. Fernald includes an adjustment for capacity utilization to make the series comparable to the theoretical benchmark. Finally, Fernald also includes an adjustment for human capital, following Mankiw et al. (1992). Formally, BCL assume that $L_t = H_t$, total hours worked, while Fernald assumes $L_t = Q_t H_t$ where Q_t is an index of labor quality based on education. Using (1), we see that the Fernald's measure comparable to the BLC measure is

$$A_t^{FQ} = A_t^F Q_t^{1-\alpha}$$

where A_t^F is Fernald's labor-quality-adjusted TFP measure. Figure 12 in the Appendix shows the three TFP series, where A_t^{BCL} is normalized to 1 in 1947 to be comparable with Fernald's measures. The key point is that none of the series is well described by the exponential process (2) with constant g. A_t^{FQ} and A_t^{BCL} are well described by the additive process (3) with constant b. The A_t^F displays some slow down in the later part of the sample even according to (3) because some of the measured productivity gains are attributed to the labor quality factor. This speaks to the model specification issue that I discuss in details in Section 3.

1.2 Postwar U.S. TFP

The simplest way to start comparing model D and model G is to consider the following experiment. Suppose that two agents, *George* and *Daniela*, are asked in the middle of the sample (1983) to predict the level of TFP in the second half of the sample (1984-2019). The agents have access to data from the end of World War 2 until 1983. The two agents have dogmatic beliefs regarding the correct model of economic growth. George believes in model G from equation (2) while Daniela believes in model D from equation (3). George therefore fits a log linear model over the years 1947 : 1983 and predicts future (log) TFP as

$$\log\left(\hat{A}_t^{(G)}\right) = \hat{a}_g + \hat{g}t$$

for $t = 1984 : 2019$. Daniela instead fits a linear model and predicts future TFP as $\hat{A}_t^{(D)} = \hat{a} + \hat{b}t$. Figure 1 shows that Daniela would have made a much better forecast than George. George is puzzled by the TFP slowdown while Daniela does not perceive an obvious long term break in her model (although there are some meaningful medium term deviations). The results obtained from A_t^{FQ} and A_t^{BCL} are virtually identical so I

Figure 1: Out-of-Sample TFP Forecasts

(a) BCL Data (b) Fernald Data

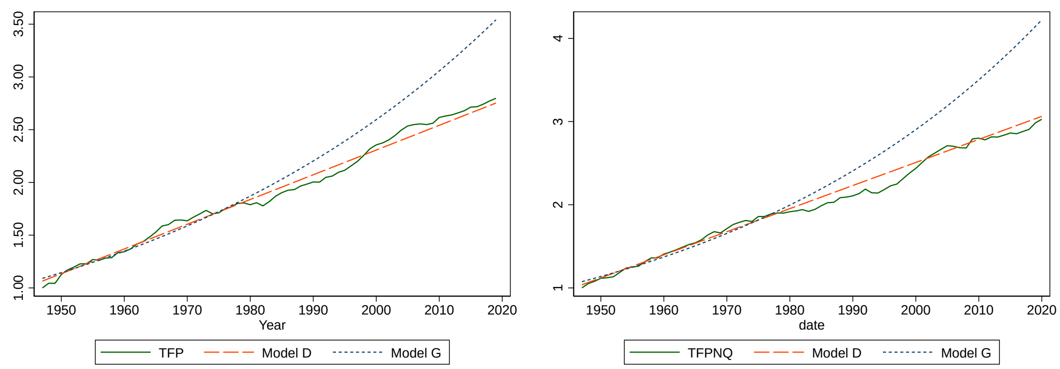

Notes: BCL TFP is in based on $US 2010. Fernald unadjusted TFP, $A_t^{FQ} = A_t^F Q_t^{1-\alpha}$. Both are normalized to 1 in 1947. Models are estimated over 1947-1983. The forecast 1984-2019 is out-of-sample. Data source: Fernald (2012) and Bergeaud et al. (2016).

focus on one measure (BCL) for brevity.

Figure 1 reveals a new fact and makes an important empirical point. The new fact is that there is no TFP slowdown in the US according to model D.

Fact 1. *There is no TFP slowdown in the US according to model D.*

The important empirical point is that one should always use out-of-sample forecasts to test growth models and that, with realistic values for TFP growth rates, distinguishing between models D and G requires at least 10 years of out-of-sample forecasts. In-sample fit always looks deceptively good: models D and G appear to fit equally well the 1947-1983 sample. Figure 13 in the Appendix shows that in-sample fit does not provide a useful diagnostic even for the most basic measure (real GDP per capita) and even over a long sample (130 years).

1.3 Capital Accumulation and Labor Productivity

Let us know 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

$$k_t^{1-\alpha} = \frac{\alpha}{\chi} A_t. \tag{4}$$

Equation (4) says that the normalized inverse MPK (IMPK) is proportional to A.[2] Model G therefore predicts that $k_t^{1-\alpha}$ grows exponentially, while model D says that it grows linearly. Figure 2 presents the forecasts for $k_t^{1-\alpha}$ based on models D and G with $\alpha = 0.3$, the value used by BCL. For model D we have

$$\mathbb{E}\left[k_t^{1-\alpha}\right] = \hat{a}_{impk} + \hat{b}_{impk} t \tag{5}$$

For model G we have the formula in logs. Once again we find that the log-linear model with constant growth widely missed the mark, while the additive model gives a useful forecast. This test is obviously a test of the joint hypothesis of linear TFP growth and a constant user cost, together with Cobb-Douglass capital demand. This last assumption is certainly not correct in many cases, but the data reveals that it may still provide a useful approximation. The main reason for fitting equation (5), however, is to be able to forecast labor productivity (and GDP per capita).

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:

$$\lambda_t \equiv \frac{Y_t}{L_t} = A_t k_t^{\alpha}. \tag{6}$$

Model D offers a forecast for labor productivity as

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

Note that labor productivity is convex in time even under additive growth since it depends on the *product* of both TFP and capital intensity. I could use the forecasts for IMPK and TFP to similarly create a forecast for model G but that would be a straw man since we have already seen that model G fails to predict either A or IMPK. To give model G a chance, I create directly a forecast of labor productivity by fitting the

[2] Users of model G typically interpret equation (4) as saying that capital grows grows exponentially, just like A, as a rate $(1+g)^{1/(1-\alpha)}$. Equivalently, if the model is written with Harrod-neutral technological progress, $Y_t = K_t^{\alpha}\left(\mathcal{Z}_t H_t\right)^{1-\alpha}$ then capital is proportional to $\mathcal{Z}_t$. I return to these issues in Section 3.

Figure 2: Out-of-Sample $IMPK$ and LP Forecasts

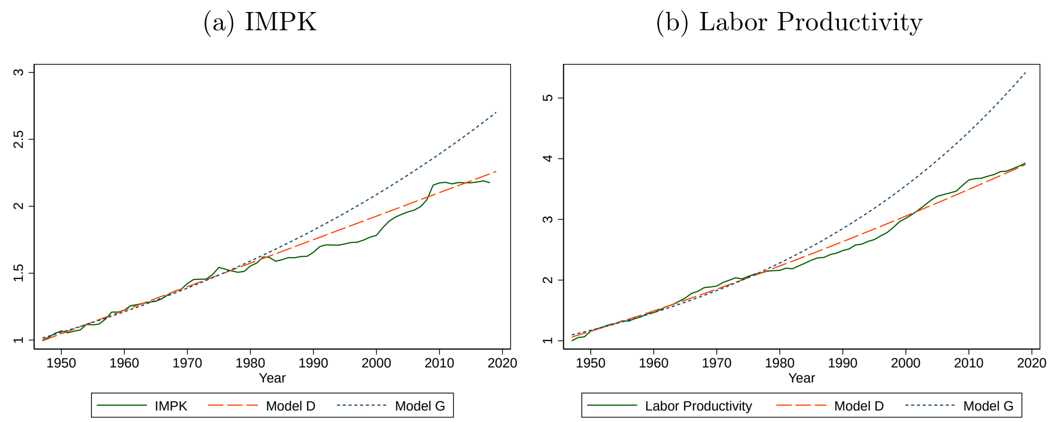

Notes: $IMPK = (K/L)^{0.7}$, normalized to 1 in 1947. Models are estimated over 1947-1983. The forecast 1984-2019 is out-of-sample. Labor productivity is real GDP per hour. Data source: Bergeaud et al. (2016).

series for $\log(\lambda_t)$ in the first half of the sample. Model G therefore gains two degrees of freedom. Panel (b) in Figure 2 shows that the convex-linear forecast of model D predicts correctly the evolution of labor productivity in the long term. Model G does not.

US growth is better described as additive rather than multiplicative. Instead of stating that the average growth rate of TFP is 1.45%, which is correct but not particularly useful, it is more relevant to say that TFP increases by 0.0245 points each year starting from a normalized value of 1 in 1947. For labor productivity, both the additive growth model and the multiplicative growth model predict an increasing size of productivity increments, but at different speeds. The additive model D predicts that labor productivity increments increase with the square of the time horizon, while the geometric model G predicts exponentially increasing increments. Model G does not describe the data with a constant growth rate. Model D describes the data relatively well with year-on-year increments of about \$1560 per full time worker (\$0.87 per hour, assuming 1800 hours worked in a year) around 2010.

Fact 2. *Postwar US TFP growth is well described by Model D with increments of* $\Delta = 0.0245$ *points each year starting from a normalized value of 1 in 1947. Model D also predicts the correct non-linear evolution of labor productivity.*

1.4 U.S. TFP, 1890-2019

Let us know extend the methodology and the sample, taking into account that trend growth can change over time. Our agents now forecast time varying growth according to a standard exponential smoothing model

$$\mathbb{E}_t[\Delta_{t+1}] = (1-\zeta)\,\mathbb{E}_{t-1}[\Delta_t] + \zeta\Delta_t \tag{7}$$

with $\Delta_t \equiv A_t - A_{t-1}$ for model D and $g_t = A_t/A_{t-1} - 1$ instead of Δ_t for model G.

There are two ways to set the smoothing parameter. One can argue on theoretical grounds that changes in the trend growth rates of TFP are decadal phenomena. This approach suggests values of ζ between 0.05 and 0.1. At 0.05, the sensitivity of the trend estimate to the most recent observation is the same as that of a 20-year moving average. Below 0.05 the model would take too long to adjust to changing trend growth. At 0.1 the sensitivity would be the same as that of a 10-year moving average. The main advantage of this approach is that it avoids any risk of over-fitting or p-hacking. I will simply report the results for 0.05 and 0.1 (and intermediate values) and see if the results are robust.

Figure 3 shows the raw and smoothed series for $\zeta = 0.05$ and $\zeta = 0.1$. The data is from Bergeaud et al. (2016) and winsorized in the first and last percentiles to remove limit extreme outliers during WW2. The model is initiated over the first 10 observations, 1891 to 1900. As expected the trend growth of the economy changes over this long sample.

The other way to choose ζ is to estimate it in some sample. The advantage is obvious, but the cost is that we waste a sample where we cannot perform out-of-sample tests. Thankfully the two approaches turn out to yield similar results. The smoothing parameter that minimizes the RMSE of one-year forecasts from (7) for the US over 1890-2019 is $\zeta = 0.0664$. If we consider the RMSE of 10-year ahead forecasts, the optimal parameter is $\zeta = 0.055$ (see below for this calculation). In the remaining of the paper I will therefore use $\zeta = 0.05$.

The extreme heteroskedasticity of growth rates is also apparent in Panel (b) of Figure 3. TFP growth rates are much more volatile before than after WW2 – 4.9% vs 1.5% – and volatility declines further after 1980. Romer (1986a) discusses the first fact, McConnell and Perez-Quiros (2000) discuss the second fact, which became known as the great moderation puzzle. These puzzles do not exist in Model D. The standard deviation of TFP changes is 0.13 before WW2 and 0.11 since 1947, and the difference is

Figure 3: US TFP, 1890-2019

(a) TFP Changes. Model D. (b) TFP Growth Rates. Model G.

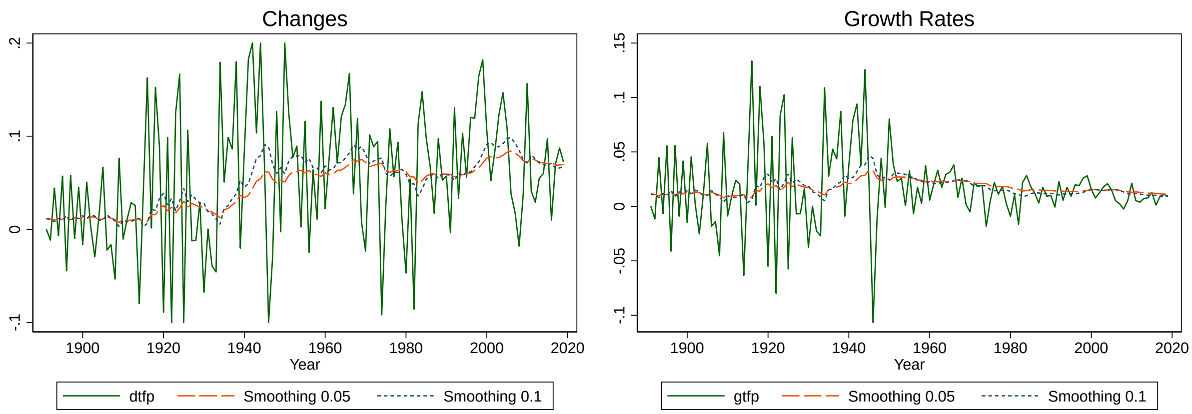

Notes: Models are estimated over 1947-1980. The left panel show the prediction of a linear model. The right panel shows the prediction of a log-linear model. US TFP is from the updated work of Bergeaud et al. (2016).

not statistically significant. Formally, define the residuals for model G as

$$\eta_t^g = g_t - \mathbb{E}_{t-1}\left[g_t\right]$$

and similarly for model D, $\eta_t^{\Delta} = \Delta_t - \mathbb{E}_{t-1}\left[\Delta_t\right]$. Table 1 shows that the volatility of TFP growth rates declines significantly over time. I use the absolute value of the unexpected shock to avoid the influence of outliers but the results are similar if I use squared residuals instead, as in ARCH models. Average absolute deviation is 2.5% in the sample, and decline by 3.6 basis point each year on average. Over 50 years the volatility changes by 1.8% which is almost 3/4 of the sample average. By contrast the trend is small and insignificant for model D. The change over 50 years is only 10% of the sample average.

Fact 3. *There is no volatility puzzle for model D.*

Forecasts Let us know study the forecasting accuracy of the two models as in Figure 1, but instead of performing once test pre/post 1980, I compute real-time rolling estimates and out-of-sample forecasts using (7). Figure 4 shows the 10-year out-of-sample predictions for the level of TFP from model D

$$\mathbb{E}_{t-10}^{D}\left[A_t\right] = A_{t-10} + 10\hat{b}_{t-10} \tag{8}$$

Table 1: Volatility of TFP Growth, US 1890-2019

100*	Model G $\lvert\eta_t^g\rvert$	Model D $\lvert\eta_t^\Delta\rvert$
(Year-1955)	-0.036	-0.011
t-stat	-6.5	-1.1
Constant	2.49	5.62
t-stat	11.8	15.1
N	129	129
R^2	0.247	0.009

Notes: Dependent variables scaled by 100. For model G the dependent variable is residual growth rate of TFP. For model D the dependent variable is the residual of the first difference of TFP. Data from Bergeaud et al. (2016), US, 1890-2019.

Table 2: RMSE for US TFP Forecasts, 1890-2019

Smoothing Parameter	$\zeta = 0.05$		$\zeta = 0.1$	
Forecast Horizon	10 years	20 years	10 years	20 years
Model D	.086	.145	.090	.147
Model G	.107	.209	.114	.237

Notes: US TFP is from the updated work of Bergeaud et al. (2016)

and from model G

$$\mathbb{E}_{t-10}^{G}\left[A_t\right] = A_{t-10}\left(1+\hat{g}_{t-10}\right)^{10}, \tag{9}$$

where $\hat{b}_{t-10} \equiv \mathbb{E}_{t-10}\left[\Delta_{t-9}\right]$, $\hat{g}_{t-10} = \mathbb{E}_{t-10}\left[g_{t-9}\right]$ are the trends estimated 10 years before and I use the fact that $\mathbb{E}_t\left[\Delta_{t+k}\right] = \mathbb{E}_t\left[\Delta_{t+1}\right] = \hat{b}_t$ for all $k \geq 1$. I then define the long term forecast error as

$$\epsilon_t^{D,G} = \frac{A_t - \mathbb{E}_{t-10}^{D,G}\left[A_{i,t}\right]}{\bar{A}},$$

where $\bar{A}$ is the sample average of A. I use this normalization to ease the comparison across datasets where TFP levels are defined in different ways. Table 2 reports the root mean square errors (RMSE) of long term forecasts. Model D outperforms model G in all cases and the relative performance of model D *increases* with the forecast horizon.[3] The main reason is that after a sequence of positive growth rates the multiplicative model extrapolates exponential growth for 10 years, which systematically fails to materialize.

Fact 4. *For US TFP over 1890-2019, model G's long-term forecast errors are 25% to 40% higher than those of Model D.*

[3]The same results hold if I compute the RMSE over relative errors $\frac{A_{i,t}-\mathbb{E}_{t-10}^{D,G}[A_{i,t}]}{A_{i,t}}$.

Figure 4: US TFP Forecasts, Long Sample

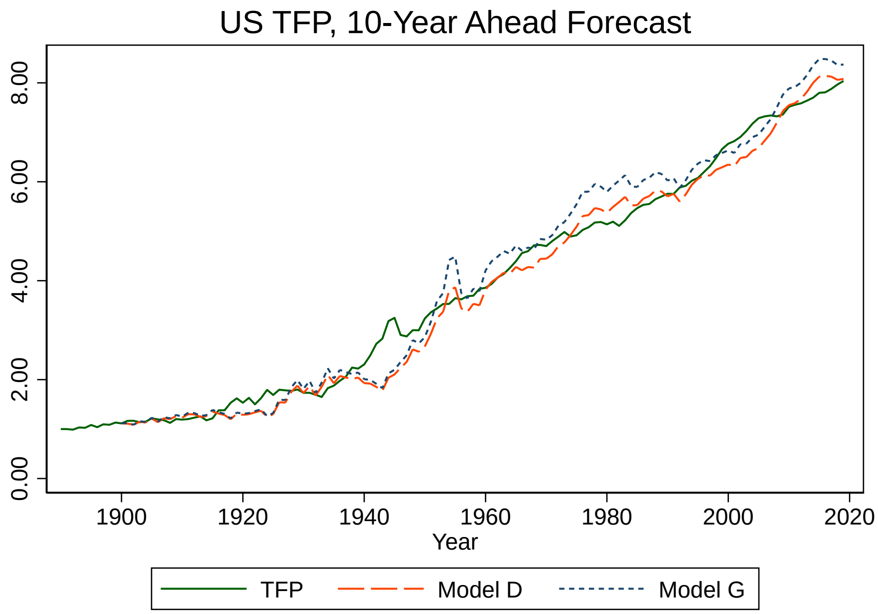

Notes: Forecast with smoothing parameter 0.05. US TFP is from the updated work of Bergeaud et al. (2016).

Figure 3 shows that model G is unstable. The estimated trend growth rate is constantly being revised. This is why model G is not useful as a long run growth model. To illustrate the point consider the predictions one would make in 2020 regarding GDP in 2060, holding population constant so as not to introduce additional demographic forecast errors. Using Fernald's data, TFP is 3 in 2020. The estimate for TFP growth is 1.2% with a standard deviation of 0.2% over the preceding 40 years. The estimate for TFP increments is 0.027 with a standard deviation of 0.0036. The G-forecast for cumulative growth between 2020 and 2060 is 2 (i.e., $1.012^{\frac{40}{1-\alpha}}$) but the two standard errors range is 1.6 to 2.6, which is one entire GDP of 2020, or \$21 trillion. It is difficult to see the usefulness of a forecast with such a wide error range. The D-forecast is 1.59 with a range of 1.42 to 1.76, which is only one third of 2020 GDP.

2 Country-Level International Evidence

Bergeaud et al. (2016) provide data for 23 countries.[4] The trend growths are estimated

[4] Australia, Austria, Belgium, Canada, Switzerland, Chile, Germany, Denmark, Spain, Finland, France, United Kingdom, Greece, Ireland, Italy, Japan, Mexico, Netherlands, Norway, New Zealand,

Table 3: RMSE for 23 Countries, BCL Sample

Sample	1890-2019		1950-2019	
Parameter	$\zeta = 0.05$	$\zeta = 0.1$	$\zeta = 0.05$	$\zeta = 0.1$
Model D	.130	.128	.102	.103
Model G	.171	.168	.162	.145
N. Obs.	23	23	23	23

Notes: Data from Bergeaud et al. (2016).

with the recursive learning model (7) with parameter $\zeta = 0.05$ and $\zeta = 0.1$. As before, all the forecasts are out-of-sample. For each country $i = 1 : 23$ and each year t I compute the forecast errors as

$$\epsilon_{i,t}^{D,G} = \frac{A_{i,t} - \mathbb{E}_{t-10}^{D,G}\left[A_{i,t}\right]}{\bar{A}_i},$$

where $\bar{A}_i$ is the country sample average and the expectation are taken under models D and G. Finally, I compute the root mean square error for each country as

$$\texttt{RMSE}_i^{D,G} = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(\epsilon_{i,t}^{D,G}\right)^2}.$$

2.1 Long-Sample, 1890-2019

I first run the model over the whole sample, 1890-2019, initializing over the first 10 years. Figure 5 shows that the D-model out-performs the G-model for every single country in the BCL sample.

Table 3 summarizes the average performance of models D and G. The differences are larger than in Table 2 because many countries experience more volatile growth sequences than the US, which makes it easier to separate the two models. Model D over-performs model G by 30% to 60%.

2.2 Post-War Sample

I run the model separately for the post war period because several countries (e.g. Japan, Germany) experience large shocks during the 1940s which may render the forecasts from

Portugal, Sweden and United States. The sample covers 1890–2019. The main variables are GDP, labor, and capital. Labor is constructed from data on total employment and working time. Capital is constructed by the perpetual inventory method applied equipment and buildings.

Figure 5: TFP Forecast Errors, BCL 1890-2019

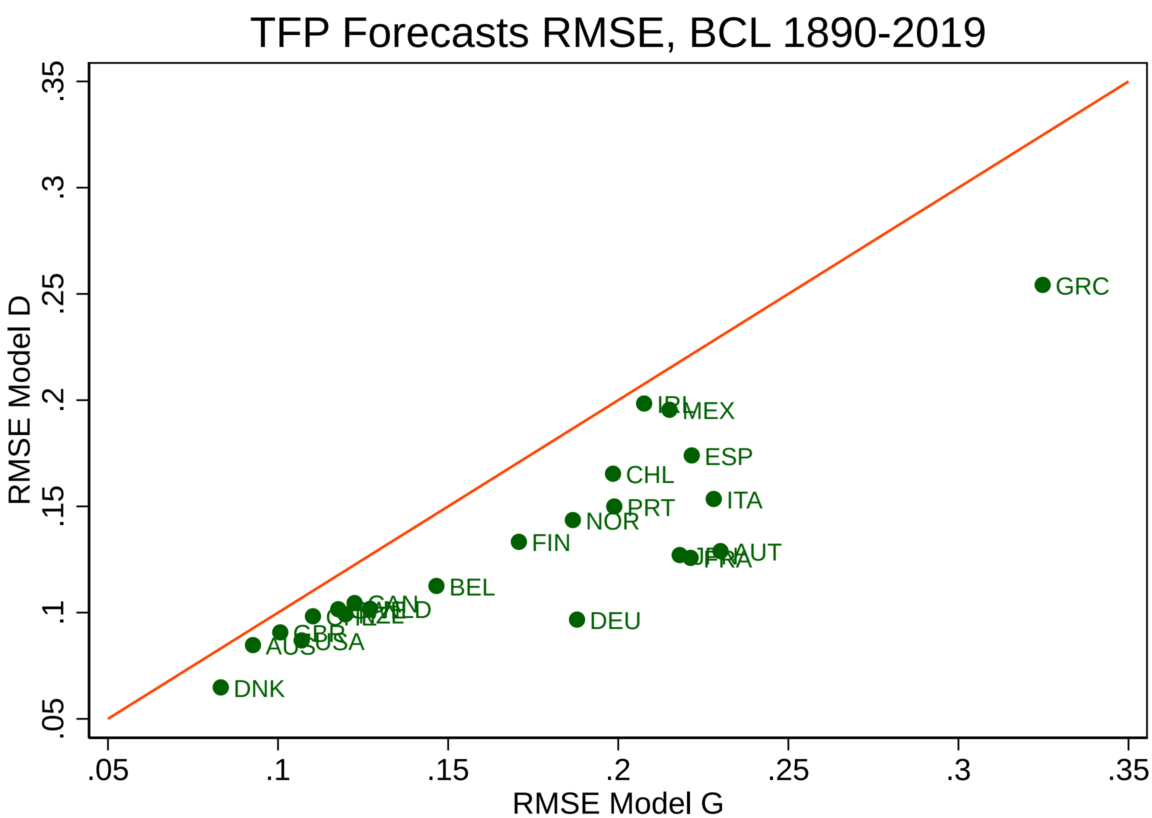

Notes: US TFP is from the updated work of Bergeaud et al. (2016).

the exponential model unstable. Figure 6(a) shows the RMSE of TFP forecasts in the two datasets. Model D performs better than model G in all cases.

I also use the OECD MFP database as a robustness check in Figure 6(b). The data covers 24 countries and starts in 1985 for most, and later for some. Because the time series are much shorter it is more difficult to tell the models apart and some countries are bunched close to the 45 degree line. Nonetheless, model G never performs better than model D, and often performs worse. Perhaps the most interesting case is that of Korea, which is not in the BCL sample and has experienced strong growth over the past 30 years. It turns out that Korean TFP growth is very linear.

The OECD data does not include some important Asian countries with strong growth performance. Figure shows TFP for Thailand and Taiwan. Taiwan's TFP growth is remarkable. The TFP index, normalized to 1 in 2017, was only 0.2 in 1955. Such a fast growth makes it easy to tell apart model D and model G. Model D fits very well. Model G vastly over-predicts TFP, irrespective the smoothing parameter. This is consistent with the argument in Young (1995) that TFP growth in East Asia was not exceptional.

Fact 5. *TFP growth is better described by model D than by model G for both developed and developing countries.*

2.3 TFP Catch-Up and Slowdown

Section 1 has shown that there is no TFP slowdown in the US once we use the correct linear benchmark. There is, however, a TFP slowdown in the euro area (EA, defined as of current membership) and in Japan. Figure 8 compares the evolution of TFP (left panel) and labor productivity (right panel) in the US and the EA during the post-war era. From 1947 to 1980 the EA catches up with the US. Between 1980 and 1990 TFP grows somewhat faster in the EA than in the US. From the mid 1990s onward, however, EA TFP starts to fall behind US TFP. The right panel shows the same pattern for labor productivity, and the dashed line shows that the entire difference in output per hour between the two regions comes from differences in TFP, not in capital accumulation.

We observe a (linear) TFP slowdown in some but not all countries, as show in Table 4. The sample is split in 1991 because this is the peak of the EA relative TFP performance. The table reports the average of $\Delta\left[TFP\right]/TFP_{US,1947}$ since this is the relevant measure under linear TFP growth. The US grows at a roughly constant increment of its 1947 TFP, and up to 1990 the EA and Japan grow at a faster pace. After 1990, however, their TFP increments decline dramatically and fall below that of the US. The key point

Figure 6: TFP Forecast Errors, Post War

(a) BCL 1950-2019

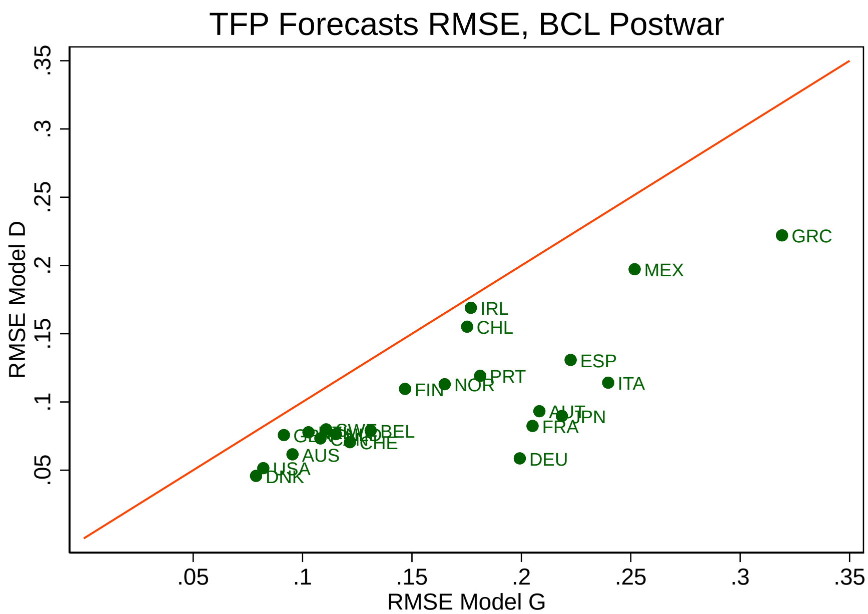

(b) OECD, post-1985

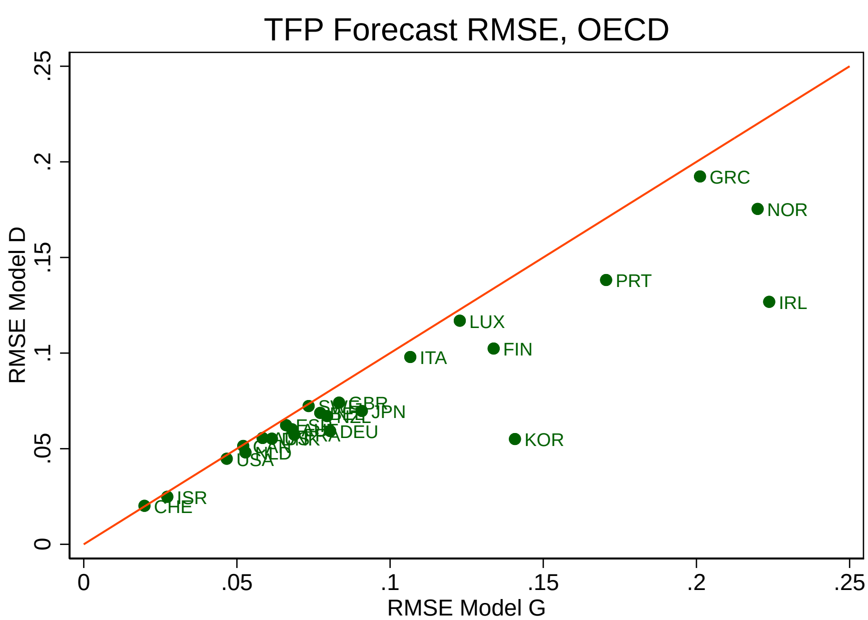

Notes: Model G on the horizontal axis, model D on the vertical axis. Out-of-sample, 10-year forecasts with smoothing parameter 0.05. Data from Bergeaud et al. (2016). Sample 1950-2019.

Figure 7: TFP, Fast Growing Asian Countries

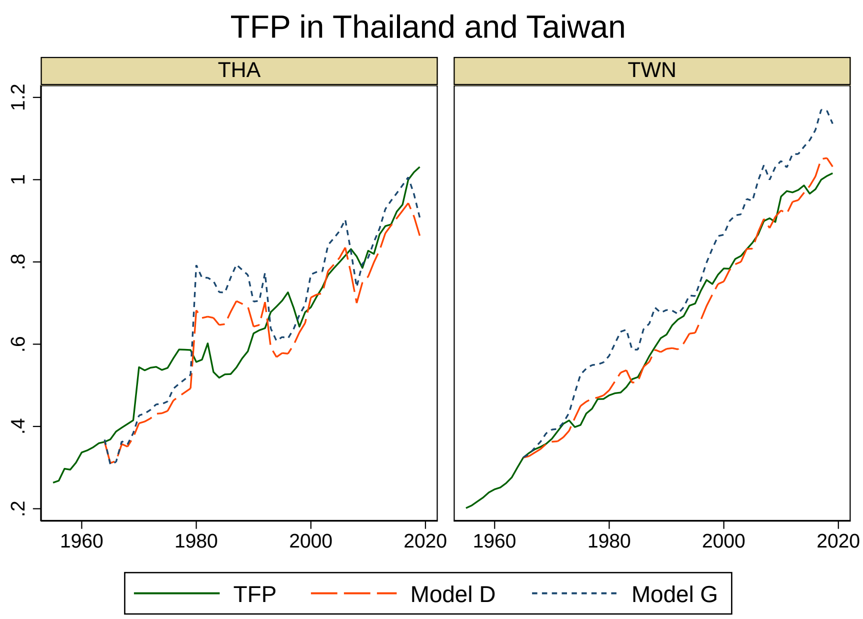

Notes: Data from Penn tables Asia.

Figure 8: TFP and Labor Productivity in the US and the Euro Area

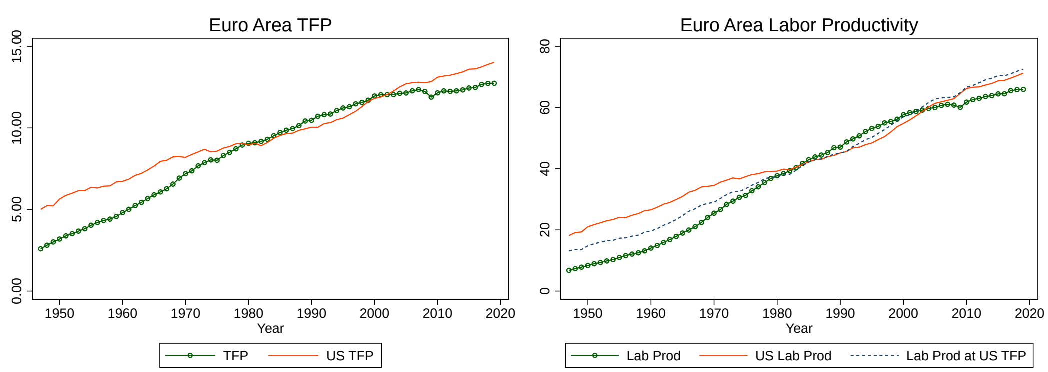

Sources: BCL. In the right panel, labor productivity is GDP per hour. Using equation (6) EA labor productivity at US TFP is defined as $\tilde{\lambda}_{EA,t} = A_{US,t} k_{EA,t}^{\alpha}$.

Table 4: TFP Increments

$\Delta\left[TFP\right]/TFP_{US,1947}$	1947-1990	1991-2019
USA	.023	.027
Euro Area	.037	.016
Japan	.029	.012
Denmark	.026	.026
Sweden	.022	.028

Notes: TFP Increments measured in units of US TFP in 1947. Data from Fernald (2012).

is that this slowdown goes beyond what one might expect at the end of the catch-up period.

All models are wrong; some are useful. The exponential model is not useful to think about long run TFP growth because no country is able to live up to its extreme predictions.[5] It leads us to conclude that the TFP slowdown is a universal phenomenon, leaving little room for policy to explain differences between countries. By contrast the linear TFP benchmark highlights that Denmark and Sweden have TFP performances comparable to that of the US while Japan and the Euro area do not. On a related note, the exponential model can paint a misleading picture of growth in countries with high TFP. The danish TFP increment, at 0.026 is much higher than that of the Euro Area at 0.016. The difference in growth rates is less impressive, however, because TFP in Denmark in 2019 is 13% higher than in the Euro area.

3 Theoretical Implications

This section highlights some important features of additive growth. To draw long-run implications from the theory we must first revisit the issue of the production function.

3.1 Finding Linearity

As explained in Barro and Sala-i-Martin (2004), balanced growth requires labour-augmenting (Harrod-neutral) technology

$$Y_t = F\left(K_t, A_t^L L_t\right) \tag{10}$$

[5] Unless the predicted growth rate is continuously revised downward so as to emulate a linear growth model, but that merely proves the point that the exponential benchmark is useless.

where A_t^L is labor-augmenting, or Harrod-neutral, technological progress. We need an exact mapping between this functional form and the evidence discussed so far even if $F(.)$ is Cobb-Douglas, as Fernald (2012) and Bergeaud et al. (2016) assume. In that case we can of course renormalize our productivity measure as $A^{\frac{1}{1-\alpha}} = A^L$ so that (10) becomes $Y_t = A_t K_t^\alpha L_t^{1-\alpha}$, but this does not answer the question of whether A or A^L (or perhaps neither) is best described as a linear process. To be concrete, if A^L is linear then A is concave in time. If instead A is linear then A^L is convex in time. These distinctions matter for long-run growth, as I discuss later.

Fernald (2012) writes a production function

$$Y_t = A_t^F K_t^\alpha \left(Q_t H_t\right)^{1-\alpha},$$

where A_t^F Fernald's headline TFP measure, capital services K_t are constructed from disaggregated series on structures, equipments and IPs and adjusted for variable utilization, while Q_t is a labor-quality index based on Mincer wage regressions from the Current Population Survey. I consider 5 hypotheses to describe US TFP:

- (i) GA: the Hicks-neutral series net of labor quality improvements, A_t^F, features constant *exponential* growth. This is the textbook assumption.
- (ii) DA: A_t^F features constant *additive* growth.
- (iii) DAQ: $A_t^{FQ} = A_t^F Q_t^{1-\alpha}$, the Hicks-neutral TFP *including* educational improvements is additive. Whether or not one should net out the effect of education when measuring TFP depends on the question at hand. Solow (1957) explains that he uses "the phrase "technical change" as a short-hand expression for any kind of shift in the production function. Thus [..] improvements in the education of the labor force, and all sorts of things will appear as "technical change"." If one follows this line of reasoning, then $A_t^{F,Q}$ is the relevant TFP concept.
- (iv) DAL: $\left(A_t^F\right)^{\frac{1}{1-\alpha}}$, the Harrod-neutral series adjusted for labor quality, is additive. If the true model is $Y_t = F\left(K_t, A_t^L Q_t H_t\right)$ with A_t^L additive, then we should find that $\left(A_t^F\right)^{\frac{1}{1-\alpha}}$ is additive.
- (v) DALQ :$\left(A_t^F\right)^{\frac{1}{1-\alpha}} Q_t$, the Harrod-neutral series including educational improvements is additive. If the true model is $Y_t = F\left(K_t, A_t^L H_t\right)$ with A_t^L additive, then we should find that $\left(A_t^F\right)^{\frac{1}{1-\alpha}} Q_t$ is additive.

Table 5: Trends in US TFP Growth

	(i) GA	(ii) DA	(iii) DAQ	(iv) DAL	(v) DALQ	(vi) BCL
$100 * \Delta[.]$	$\log A_t$	A_t	$A_t Q_t^{1-\alpha}$	$A_t^{\frac{1}{1-\alpha}}$	$A_t^{\frac{1}{1-\alpha}} Q_t$	A_t^{BCL}
(Year-1983)	-0.026	-0.015	-0.002	-0.005	0.038	-0.003
t-stat	-3.3	-1.1	-0.1	-0.2	1.1	-0.3
Constant	1.311	2.107	2.758	4.131	5.740	2.498
t-stat	7.9	7.2	8.5	6.7	7.9	9.9
N	72	72	72	72	72	72
R^2	0.133	0.017	0.000	0.000	0.017	0.001

Notes: US, 1947-2019. A_t refers to Fernald's measures of TFP, adjusted for labor quality Q, and normalized to 1 in 1947. In column (iii) the labor quality adjustment is added back to the TFP measure. All dependent variables scaled by 100. In (i) the dependent variable is the growth rate of TFP. In (ii-vi) the dependent variable is the first difference of TFP. Data from Fernald (2012) in (i-v) and Bergeaud et al. (2016) in (vi). The independent variable is year minus sample mean (1983) so that the constant can be readily interpreted as the sample average growth in percent for (i) and in percentage points of the 1947 value in the other columns.

Table 5 shows the estimates for the US from Fernald's data, together with one estimate from the BCL data for comparison. Column (i) documents the well-know TFP "slow-down" which is a puzzle for the exponential growth model (the puzzle is the same for all G models irrespective of the labor quality or Harrod-neutral adjustments, omitted for brevity). The sample average TFP growth is 1.3% per year, but loses 2.6 basis point each year, from around 2% in the early 1950s down to only 0.5% after 2010. Column (ii) shows that the puzzle is much reduced, but not entirely eliminated, by the DA specification: if one takes away the contribution of human capital to productivity growth, the resulting TFP series is sub-linear.

Column (iii) and (iv) show that DAQ and DAL are two equally plausible way to characterize additive growth. DAQ folds educational improvements into Hicks-TFP growth instead of netting them out. DAL assumes linear labor-augmenting productivity applied to quality adjusted labor $Q_t H_t$. Column (v) shows that doing both adjustments simultaneously (DALQ) might be excessive. Column (vi) shows that the BCL series is similar to the Hicks-neutral series based on raw labor in (iii), which is consistent with our discussion in Section 1.

Table 5 suggests that models (iii) and (iv) are approximately additive. An important point is that both models predict increasing improvements in labor productivity, but for slightly different reasons. Model DAQ generates Hicks-neutral additive growth from the combination of educational and other improvements. In model DAL, A_t^L is linear, which by itself generates linear labor productivity, but it applies to an increasingly qualified quantity of labor. There is no clear statistical reason to prefer the Hicks-additive DAQ to DAL, but DAQ has two practical advantages. It is applicable to the BCL series (vi) and it requires the forecast of only one factor (A_t) instead of two (A_t^L and Q_t).

Fact 6. *Hicks-TFP A_t – including educational improvements and normalized to 1 in 1947 – for the US private sector increases linearly by about 2.76 percentage points each year.*

After twenty years, TFP is 1.55, after forty years it is 2.1, and so on. This model accounts well for the evolution of TFP in the US since 1947. There is no slowdown in TFP increments. The point estimates of -0.002 is rather precisely estimated at 0.

3.2 Theoretical Properties

Let us now turn to the theoretical dynamics of the Hicks-additive model. I use continuous time to simplify the notations. Output is $Y_t = A_t K_t^\alpha H_t^{1-\alpha}$, TFP grows according to $\frac{dA_t}{dt} = b$, and capital accumulates as

$$\frac{dK_t}{dt} = I_t - \delta K_t.$$

Hours grow at the constant population growth rate g_n: $\frac{dH_t}{dt} = g_n H_t$.

Solow-Swan Dynamics Let us start with a fixed saving rate s: $I_t = sY_t$. Define $A_t^L = A_t^{\frac{1}{1-\alpha}}$ and the scaled capital stock as

$$\kappa_t \equiv \frac{K_t}{A_t^L H_t}$$

to obtain the usual differential equation

$$\dot{\kappa}_t = s\kappa_t^\alpha - \left(\delta + g_n + \frac{\dot{A}_t^L}{A_t^L}\right)\kappa_t \tag{11}$$

with $\frac{\dot{A}_t^L}{A_t^L} = \frac{1}{1-\alpha}\frac{\dot{A}_t}{A_t}$. The exponential growth model predicts that $\frac{\dot{A}_t}{A_t}$ is constant. The additive growth model predicts that $\frac{\dot{A}_t}{A_t} = \frac{b}{A_t}$ declines over time. Since $\lim_{t\to\infty} A_t = \infty$ we have the following proposition.

Proposition 1. *The long-run additive growth path is characterized by* $\kappa_\infty = \left(\frac{s}{\delta+g_n}\right)^{\frac{1}{1-\alpha}}$. *The capital labor ratio grows as* $k_t = \kappa_\infty A_t^{\frac{1}{1-\alpha}}$ *and labor productivity as* $\lambda_t = \kappa_\infty^\alpha A_t^{\frac{1}{1-\alpha}}$. *The increments of* k_t *and* λ_t *increase to infinity,* $\lim_{t\to\infty} \frac{d\lambda_t}{dt} = \infty$, *but their growth rates converge to zero* $\lim_{t\to\infty} \frac{d\log\lambda_t}{dt} = 0$.

Let me now discuss a few implications and extensions of the model.

Long Term Growth Note that for large t we have $\frac{d\lambda_t}{dt} \approx \frac{1}{1-\alpha}\kappa_\infty^\alpha (A_0 + bt)^{\frac{\alpha}{1-\alpha}} b$. This shows that labor productivity, and thus living standards, grow as an increasing pace when we assume Hicks-linear growth. If we instead assume (counter-factually as discussed above) Harrod-linear growth, $\dot{A}_t^L = b$, then increments in living standards would not go to infinity but instead converge to a finite limit: $\frac{d\lambda_t}{dt} \to \kappa_\infty^\alpha b$. Models DAQ (iii) and DALQ (v) therefore make different predictions about long run growth. The good news is that the available evidence supports model DAQ. One caveat, however, is that model DAQ is linear because we include educational achievements into TFP growth. If the educational achievements of the 20th century as documented by Goldin and Katz (2008) cannot be repeated, growth could fall to that implied by the DALQ model. Similarly, Hsieh et al. (2019) show that human capital misallocations have decreased over the past 60 years. They argue that "*a substantial pool of innately talented women and black men in 1960 were not pursuing their comparative advantage*," and that the improved allocation of talent can explain 20 to 40 percent of labor productivity growth. If this improvement cannot be repeated our estimate of long run b could be biased upward.

Convergence Transitional dynamics are essentially the same as in the standard model, so the results on conditional convergence discussed in Barro and Sala-i-Martin (2004) are unchanged. Unconditional convergence depends on how b's vary across countries and over time. Permanent differences in b predict infinitely increasing inequality even though growth rates converge to zero in all countries. Changes in b over time predict time-varying catch-up as we saw in Section 2.3.

Neoclassical Production Function The results in Proposition 1 generalize beyond the Cobb-Douglas case. Define A_t^L as Harrod TFP in a neoclassical production function

$Y_t = F\left(K_t, A_t^L H_t\right)$. The dynamic equation becomes

$$\dot{\kappa}_t = sf\left(\kappa_t\right) - \left(\delta + g_n + \frac{\dot{A}_t^L}{A_t^L}\right)\kappa_t$$

where $f\left(\kappa\right) \equiv F\left(\kappa, 1\right)$. As before the limit solves $sf\left(\kappa_\infty\right) = \left(\delta + g_n\right)\kappa_\infty$. Define

$$\alpha_\infty \equiv \lim_{\kappa \to \kappa_\infty} \alpha\left(\kappa\right)$$

where $\alpha\left(\kappa\right) \equiv \frac{KF_K}{F}$ estimated at the point $\kappa = \frac{K}{A^L H}$. Thus α_∞ is simply the capital elasticity (or capital share) estimated at κ_∞. This model is consistent with the evidence on additive Hicks-neutral productivity growth if and only if the Harrod-neutral process is of the form $A_t^L \approx A_t^{\frac{1}{1-\alpha_\infty}}$ where A_t is additive.

Ramsey Model and Interest Rates Let me finally discuss the case where the saving rate is endogenous. Infinitely-lived households have CRRA preferences and fixed labor supply. The equilibrium is pinned down by capital accumulation and the households' Euler equation (and the usual transversality condition, omitted for brevity):

$$\dot{\kappa}_t = f\left(\kappa_t\right) - \hat{c}_t - \left(\delta + g_n + \frac{\dot{A}_t^L}{A_t^L}\right)\kappa_t,$$

$$\frac{\dot{\hat{c}}_t}{\hat{c}_t} = \sigma\left(f'\left(\kappa_t\right) - \delta - \rho\right) - \frac{\dot{A}_t^L}{A_t^L},$$

where $\hat{c}_t = \frac{C_t}{H_t A_t^L}$ is normalized consumption per capita, σ is the EIS and ρ the rate of time preference. As before, we have $\lim_{t\to\infty} \frac{\dot{A}_t^L}{A_t^L} = 0$ so the long-term balanced growth path is given by

$$f'\left(\kappa_\infty\right) = \delta + \rho$$

and

$$\hat{c}_\infty = f\left(\kappa_\infty\right) - \left(\delta + g_n\right)\kappa_\infty$$

All per capital variables grow with A_t^L. For instance, long run per capita consumption is $c_t = \hat{c}_\infty A_t^L$. The model features decreasing growth rates therefore, assuming CRRA preferences, the risk free rate falls over time and eventually converges to ρ.

3.3 Micro-foundations and Endogenous Growth

The additive model can be cast as a semi-endogenous growth model. People are employed in production L or in research R and the labor resource constraint is $R + L = N$. For simplicity, as in Jones (2021a), I ignore capital accumulation – output is given by $Y = AL$ – and I assume an exogenous labor allocation $R = \kappa N$ for some constant $\kappa < 1$ (it is straightforward to endogenize κ by equating wages and private returns to innovation). Output per capita is $y = \frac{Y}{N} = (1 - \kappa) A$.

The growth process then depends on the semi-endogenous growth equation. We can write an additive growth process as

$$\frac{dA}{dt} = \Gamma (R) . \tag{12}$$

Alternatively, if we specify $\frac{dA}{dt} = A_t \Gamma (R)$ we obtain exponential growth. The main difference between the two models lies in the strength of inter-temporal spillovers. Consider a one time deviation from a constant $\Gamma (R)$ so that research output increases by ϵ from t_0 to $t_0 + \Delta$. At any point $t > t_0 + \Delta$ TFP becomes $A_t = A_{t_0} + \Gamma (R) (t - t_0) + \epsilon \Delta$ under additive growth, and $A_t = A_{t_0} e^{\Gamma(R)(t-t_0)+\epsilon\Delta}$ under exponential growth. In other words

$$\frac{\partial A_t}{\partial \{\epsilon \Delta\}} \mid add. = 1$$

while

$$\frac{\partial A_t}{\partial \{\epsilon \Delta\}} \mid exp. = A_{t_0} e^{\Gamma(R)(t-t_0)}$$

In the exponential model, the impact on future productivity of a small change at t_0 becomes infinitely large as we extend the time horizon. Consider the following thought experiment. Suppose that the US TFP boom of the 1930's happened during the 1910's. Using our estimates (see also the following section) of a change in local growth rate from 1% to 3.3% during that window, the exponential model says that TFP today would be 58% higher and labor productivity would be twice as high. Believing in exponential growth means believing that if we had invented/implemented electricity 20 years earlier than we actually did, US GDP today would be higher by more than 20 trillion dollars (assuming constant hours per capita). This seems widely implausible. The linear model, by contrast, says that the level of TFP would be higher by 0.8 points out of baseline of 8, so 10% higher, and GDP would be about 15% higher, as opposed to 100% higher under exponential growth.

Academics often informally motivate exponential growth by arguing that an invention today increases the likelihood of future inventions. A common example is the invention of a microscope that allows more effective medical research in subsequent years. This, however, is not an example of exponential growth. If the process is simply {microscope - new molecules - better health outcomes} we have a limited inter-temporal spillover, not an exponential model. Exponential growth requires that the newly-discovered molecules feed back directly into the construction of future microscopes, and so on. An example that could fit the exponential story is that of Moore's law, i.e., the fact that the number of transistors packed into CPUs doubles (roughly) every two years. Bloom et al. (2020), however, show that the amount of resources required to achieve this outcome may have increased by a factor of 18 since the 1970s. Since the doubling time has remained constant they conclude that research productivity has decreased by a factor of 18. Guzey et al. (2021) point out that this conclusion depends entirely on the exponential benchmark. Using a linear benchmark they conclude instead that the number of transistor added per year per researcher has increased a lot. Moore's law, then, is probably better than linear but worse than exponential. Once we step out of computer manufacturing it is difficult to imagine examples that fit the exponential model. In the service sector it seems all but impossible short of a true AI revolution.

The last issue I want to discuss briefly is the population scale effect, either in level or in growth rates (Jones, 1995; Young, 1998). Population growth can overturn the additive growth prediction of equation (12). If $R_t = \kappa N_t$ grows over time then equation (12) says that $\frac{dA}{dt}$ will not be constant. This problem is the same in all growth models and not specific to the additive model. Young (1998) proposes a solution that would work in the additive model. He shows that the population scale effect can show up in the number of varieties (and the level of utility) but not in the growth rate of the economy. One can also restore additive growth by assuming strongly decreasing returns to idea production.[6] The main point is that the population scale effect is somewhat orthogonal to the exponential-versus-additive issue which, as explained above, is really about the strength of inter-temporal spillovers.

[6] Define knowledge as $\mathcal{K}$ with $\frac{d\mathcal{K}}{dt} = \gamma R$. If $A = \log(\mathcal{K})$, then $dA = \frac{d\mathcal{K}}{\mathcal{K}} = g_n$ is constant. Why would this be the case? Jones (2021b) provides a possible micro-foundation based on combinatorial growth. Suppose ideas are drawn randomly and only the best idea matters. If the number of draws grows exponentially (e.g. because of growth in the number of researchers) and if we draw from an exponential distribution, then the maximum draw grows linearly over time.

4 Historical Breaks

In this section I provide historical evidence of changes in TFP growth in the very long run. In doing so I depart from the pure statistical framework of out-of-sample forecasts as the historical discussion is easier to frame in terms of regimes separated by breaks.

4.1 US 1890-2019: TFP and GPT

Figure 3 shows that model D, unlike model G, appears to have only one break over the period 1890-2019 in the US. We can formally test this idea following Bai and Perron (2003). The unconstrained test finds one break in the $\Delta [TFP]$ series around 1930 (the point estimate is 1933). We can test H0: no breaks versus H1: break in 1933. The W statistic is 21.72 and the p-value is 0.0. I emphasize, however, that while the existence of a break is clear, the date is really an interval between the late 1920s and the beginning of WW2.

The date of the break 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.*" This period corresponds to the large scale implementation of the discoveries of the second industrial revolution: electric light, electric power, and the internal combustion engine, as discussed in Jovanovic and Rousseau (2005). Gordon (2016) points out that it is somewhat surprising that *"much of the progress occurred between 1928 and 1950*," several decades after the discoveries were made. Following David (1990), he explains the paradox by showing that the 1930s were a period of follow-on inventions, such as the perfection of the piston power-powered aircraft and television, and the increasing quality of machinery made possible by the large increases in available horsepowers and kilowatt-hours of electricity.

Following these historical insights, figure 9 proposes an interpretation of US TFP from 1890 to 2019, using linear growth with one structural break in 1933 after the electrification revolution.

We can summarize this idea in the following remark, keeping in mind that I have normalized US TFP to 1 in 1890.

Fact 7. *From 1890 to 1933, TFP increases by .017 each year until it reaches a level around 1.75 in the early 1930s. From 1933 to 2019 TFP increases by .057 each year (3.3 p.p. of its level in 1933) to reach a level around 8 at the end of the sample.*

Proposition 2 summarizes our results on US growth from 1890 to 2019.

Figure 9: US TFP under Electrification GPT Interpretation

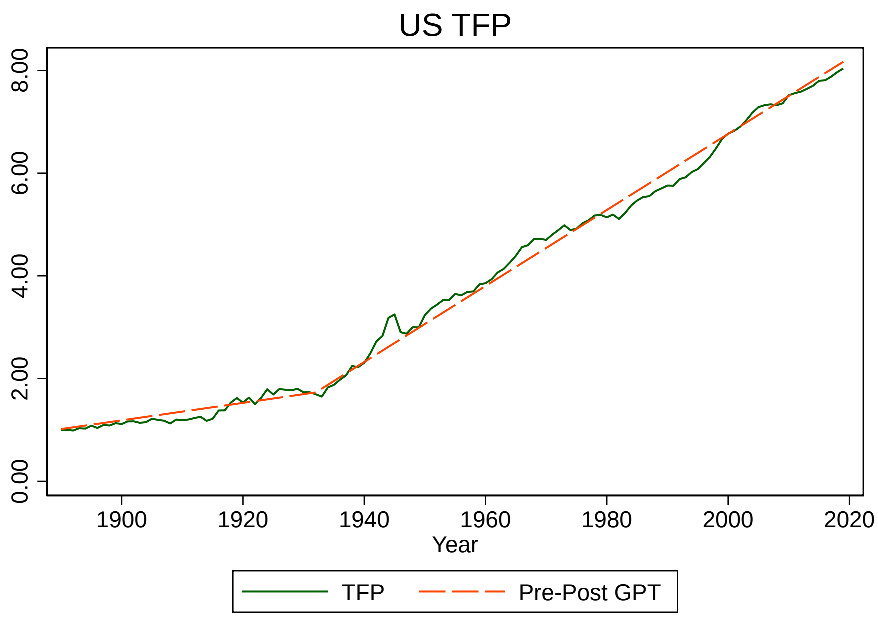

Notes: US TFP is from the updated work of Bergeaud et al. (2016), normalized to 1 in 1890.

Proposition 2. *Model G does not provide a good description of US TFP growth over 1890-2019, neither for volatility nor for long term forecasts. Model D provides a simple and accurate description as*

$$A_t - A_{t-1} = b_T + \epsilon_t$$

where $A_{1890} = 1$, ϵ_t *has a mean absolute deviation of* 0.056, $b_{1890-1933} = 0.017$ *and* $b_{1933+} = 0.057$.

It is useful to emphasize at this point that backcasting is the not the same as forecasting. My results show that the D-model offers better forecasts than the G-model, at least over a few decades. But the GPT model does not imply linear backcasts, because conditional on high productivity today, we know there must have been a break in the not-too-distant past. Thus the model does not predict that TFP was zero in 1831 (1/0.017=59 years from 1890). Instead it says that there must have been a break sometime in the 19th century. The next section shows that the data is consistent with this prediction.

4.2 UK 1600-1914

Just as the US provides the best proxy for the technological frontier in the 20th and 21st centuries, the UK provides the best proxy before 1914. I therefore perform most of the analysis from 1600 to 1914 using UK GDP per capita estimates from the Maddison Project (Bolt and van Zanden, 2020). The Maddison series for UK GDP per capita has one observation in the year 1000 and then offers annual values from 1252 onward but growth is virtually null until the 1600's.

In the neoclassical growth model, labor productivity is proportional to $A_t^{\frac{1}{1-\alpha}}$. If hours worked per capita are stationary and if the capital share is constant then we can use series on GDP per capita to construct proxies for TFP. I make these heroic assumptions and use as my proxy for TFP $(y_t)^{1-\alpha}$ where y_t is GDP per capita and $\alpha = 1/3$.

Panel (a) of Figure 10 shows the series for pseudo-TFP in the UK together with the forecasts from models D and G. TFP is normalized to 1 in 1890 so that the values are consistent with those in the BCL sample analyzed earlier. Because growth is rather slow in the 1600's and 1700's the RMSEs of the two models make rather similar forecasts. Panel (b) of Figure 10 shows the RMSEs for all the countries in the Maddison Sample over the period 1600-1914 (only a few countries have data going back to 1600, many start in the 19th century). The RMSEs of the two models are rather similar in many cases because linear and exponential forecasts are not too different when growth is slow. Australia and New Zealand provide interesting counter examples as they experience rapid growth and the G model performs very poorly, just as with Taiwan after 1960.

Figure 11 shows UK TFP together with estimated breaks of the linear model. I emphasize that these breaks are estimated using the full sample. They are *not* real time estimates of changes in trend growth as in Figure 10. Given available data, statistical agents in 1830 would not have understood the break until 1850 as we see in Figure 10. The breaks in 11 are useful to us *today* as we seek to organize the historical evidence.

Panel (b) zooms in on the two main sub-period, 1650-1830 and 1830-1914. Growth is zero until 1650 and the level of TFP is 0.4. Starting in 1650 it increases by 14 basis points each year until 1830 where it reaches approximately 0.7. In 1830 the increment increases to 58 basis point and grows linearly until WW1.

The break in 1830 is exactly as expected, but the break in 1650 happens before the first industrial revolution. There are several explanations for the fact that growth in the UK started earlier than the 18th century. The first key point to keep in mind is that I do not have a measure of hours worked. The pseudo-TFP series are based on income per-

Figure 10: Pseudo-TFP, 1600-1914

(a) UK TFP Forecasts

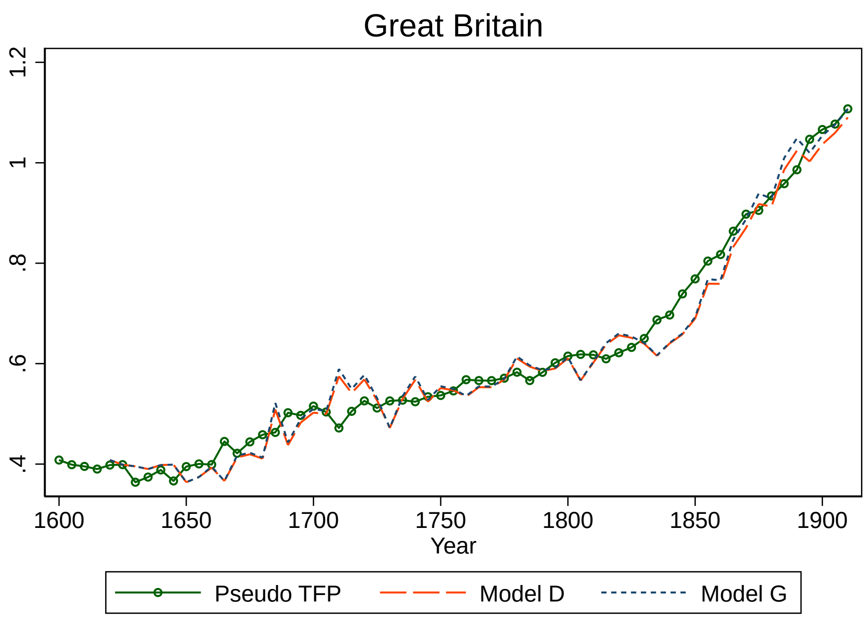

(b) RMSE

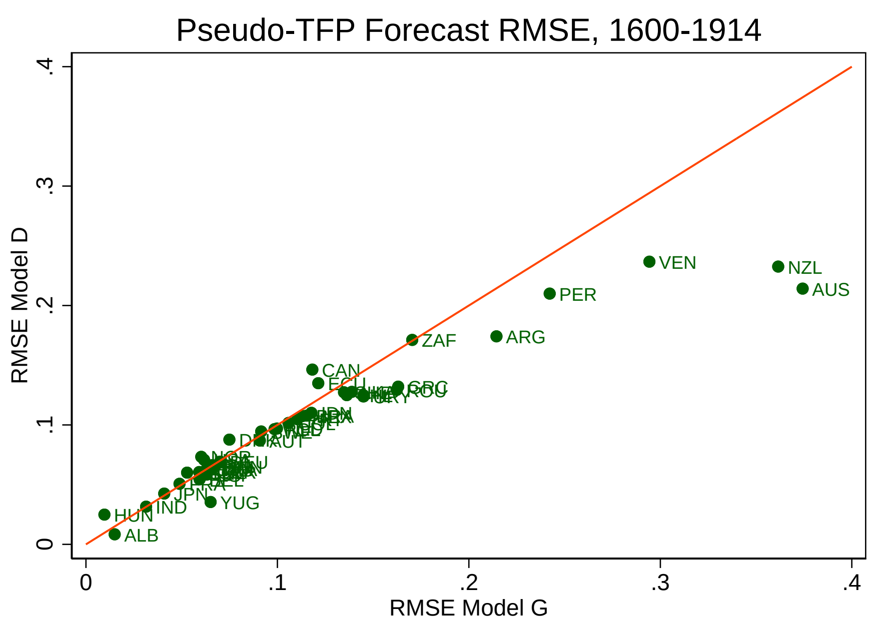

Notes: data from Maddison Project. Each circle is an average of 5 years. Pseudo-TFP is $(y_t)^{1-\alpha}$ where y_t is GDP per capita and $\alpha = 1/3$.

capita. Voth (2001) has shown that a rising labor input was an important contributor to growth after 1770. It is plausible that changes in hours per capita also contributed to growth during the previous century. Mokyr and Voth (2010) point out that "*the rise of cottage industries in the countryside after 1650, the famed "proto-industrialization" phenomenon, would do exactly that. There is also reasonable evidence to believe that labor participation rates were rising in the century before the Industrial Revolution.*" Moreover, England, unlike France, had no food crises between 1650 and 1725. Finally, the increase in GPP per capita in the 1600's is consistent with recent work by Bouscasse et al. (2021).

4.3 Rejoinder

The trend growth of the technology frontier changes enormously over time. These changes are not well predicted by the exponential model, but they do create convexity in the TFP series and they highlight the role of technological revolutions. These facts motivate a hybrid model of TFP growth. Within a GPT era, growth is linear

$$A_t - A_{t-1} = b_t,$$

but there is a small probability p of a regime change

$$b_{t+1} = \begin{cases} b_t & , 1-p \\ A_t \xi_{t+1} & , p \end{cases}$$

I normalize the new regime by the level of TFP at the time of the regime change so that the specification nests models D and G:

$$\mathbb{E}\left[A_{t+1} - A_t\right] = (1-p)\, b_t + p\bar{\xi} A_t. \tag{13}$$

Model D corresponds to $p = 0$, model G to $p = 1$. The historical data suggests $p \sim 0.5\%$ to 1% per annum which explains the success of model D. With the normalization by A_t we have $\xi_{1650} = 0.35\%$, $\xi_{1830} = 0.82\%$, and $\xi_{1930} = 3.26\%$. The structural change of the 1930s appears truly amazing in that respect.

Equation (13) is related to equation (2) in Comin et al. (2010). They point out that the exponential nature of growth depends on the complementarity between old and new technologies. In most times and places new technologies increase TFP additively. One

Figure 11: UK Pseudo-TFP & Industrial Revolutions

(a) UK TFP, Breaks

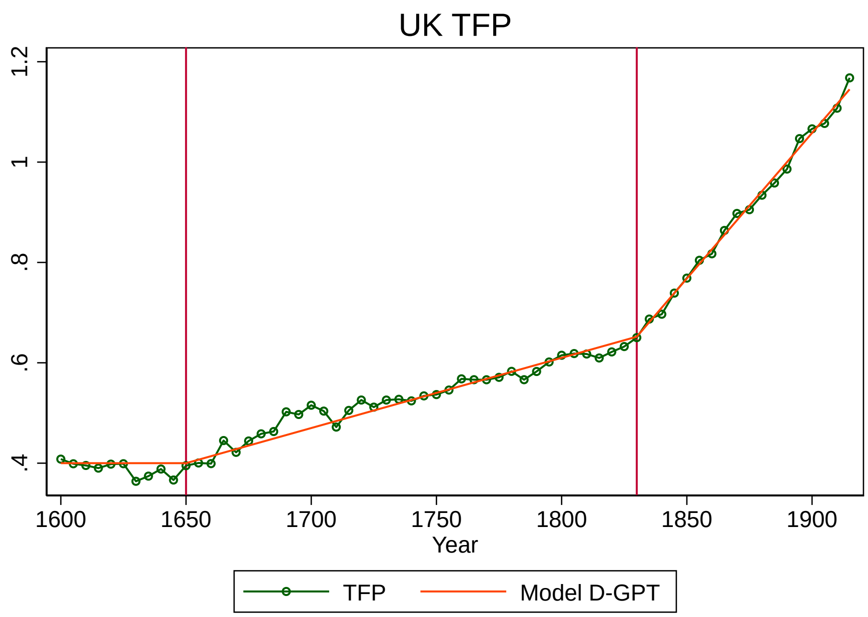

(b) UK Industrial Revolutions

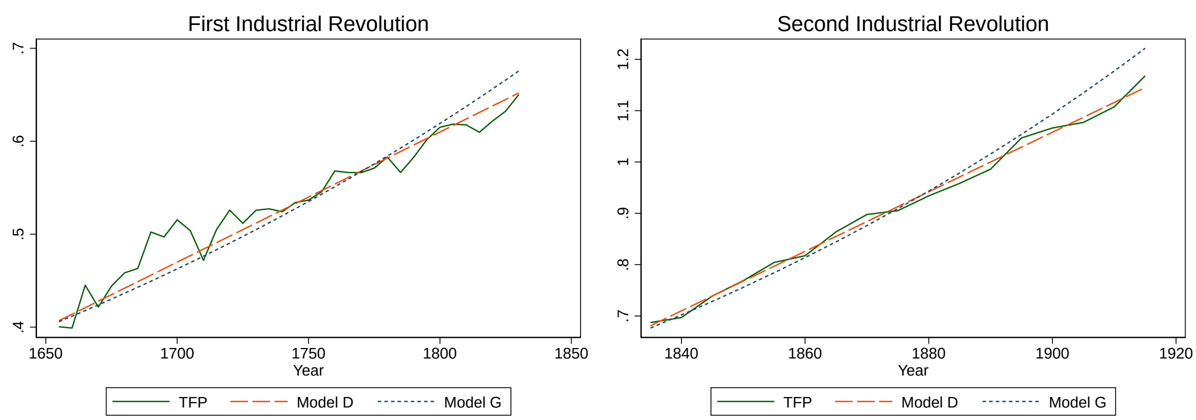

Notes: data from Maddison Project. Each circle is an average of 5 years. Pseudo-TFP is $(y_t)^{1-\alpha}$ where y_t is GDP per capita and $\alpha = 1/3$.

could interpret a GPT as a technological change that is complementary to a sufficiently high share of existing ideas and technologies. This complementarity creates what looks like multiplicative growth as the GPT is implemented.

The key point is that the TFP equation changes following the discovery of a new GPT. The linear growth equation holds within each GPT era but not across GPTs. An important question for future research is the persistence of GPTs. Should we assume that a GPT permanently increases the (potential) growth of the economy? Or should we assume that the impact on b depreciates over time? One could speculate that the slowdown of the late 1970s in Figure 9 reflects the waning impact of the initial electricity revolution and the pickup in the late 1980s the impact of IT. This is an interesting issue for future research.

5 Conclusion

TFP growth is not exponential. New ideas add to our stock of knowledge; they do not multiply it. TFP has been growing linearly over the past 90 years in the US and the additive model beats the exponential model for every single country – developed or catching up – where TFP data is available. The TFP frontier appears to grow linearly within broad historical periods: 1650 to 1830, 1830 to 1930, and 1930 until today. Additive TFP growth predicts *increasing* growth of labor productivity and GDP per capita thanks to capital accumulation. This prediction also appears to be empirically accurate.

The additive growth model explains the observed TFP slowdown as a simple side effect of model misspecification. We should not have expected growth *rates* to be constant in the first place. The additive model does not necessarily solve the research productivity puzzle of Bloom et al. (2017) since this puzzle is not about the stochastic process for TFP but rather about the specification of the production function for ideas. Models where ideas are non-rival often imply a tight connection between growth and the quantity of research. These models predict accelerating growth – whether of the linear kind or not – from an increasing number of researchers.

Additive growth has implications for industry dynamics and structural transformation as well as for firms dynamics. Philippon (2022) provides some early evidence on these issues. Additive growth may also have important implications for investment and the valuation of long term assets, such as stocks and pensions. The model predicts falling growth rates and falling interest rates so valuation and investment dynamics will depend on preferences.

References

Aghion, P. and P. Howitt (1992): "A Model of Growth Through Creative Destruction," *Econometrica*, 60, 323–351.

Bai, J. and P. Perron (2003): "Computation and analysis of multiple structural change models," *Journal of Applied Econometrics*, 18, 1–22.

Barro, R. J. and X. Sala-i-Martin (2004): *Economic Growth*, Cambridge, Massachusetts: MIT Press.

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. V. Reenen, and M. Webb (2017): "Are Ideas Getting Harder to Find?" NBER Working Paper 23782.

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

Bolt, J. and J. L. van Zanden (2020): "Maddison style estimates of the evolution of the world economy. A new 2020 update," .

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

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.

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.

Goldin, C. and L. F. Katz (2008): *The Race Between Education and Technology*, Cambridge: Belknap Harvard University Press.

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/.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019): "The Allocation of Talent and U.S. Economic Growth," *Econometrica*, 87, 1439–1474.

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.

——— (2021a): "The Past and Future of Economic Growth: A Semi-Endogenous Perspective," Annual Review of Economics.

——— (2021b): "Recipes and Economic Growth: A Combinatorial March Down an Exponential Tail," Working Paper.

Jovanovic, B. and P. L. Rousseau (2005): "General Purpose Technologies," in *Handbook of Economic Growth*, ed. by P. Aghion and S. Durlauf, Elsevier, chapter 18.

Lucas, R. E. (1988): "On the Mechanics of Economic Development," *Journal of Monetary Economics*, 22, 3–42.

Mankiw, N. G., D. Romer, and D. N. Weil (1992): "A Contribution to the Empirics of Economic Growth*," *The Quarterly Journal of Economics*, 107, 407–437.

McConnell, M. and G. Perez-Quiros (2000): "Output Fluctuations in the United States: What has Changed Since the Early 1980's," *American Economic Review*, 90, 1464–1476.

Mokyr, J. and H.-J. Voth (2010): "Understanding growth in Europe, 1700-1870: theory and evidence," in *The Cambridge Economic History of Modern Europe*, ed. by S. Broadberry and K. H. O'Rourke, Cambridge University Press, vol. 1.

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

Romer, C. (1986a): "Spurious Volatility in Historical Unemployment Data," *Journal of Political Economy*, 94, 1–37.

Romer, P. M. (1986b): "Increasing Returns and Long Run Growth," *Journal of Political Economy*, 94, 1002–1037.

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.

Voth, H.-J. (2001): *Time and Work in England, 1750-1830*, Oxford University Press.

Young, A. (1995): "The Tyranny of Numbers: Confronting the Statistical Realities of the East Asian Growth Experience," *The Quarterly Journal of Economics*, 110, 641–680.

——— (1998): "Growth without Scale Effects," *Journal of Political Economy*, 106, 41–63.

Appendix

A Three Measures of Post-War US TFP

Figure 12 compares the TFP series from BCL and Fernald, with and without adjustment for education.

Figure 12: US TFP Levels

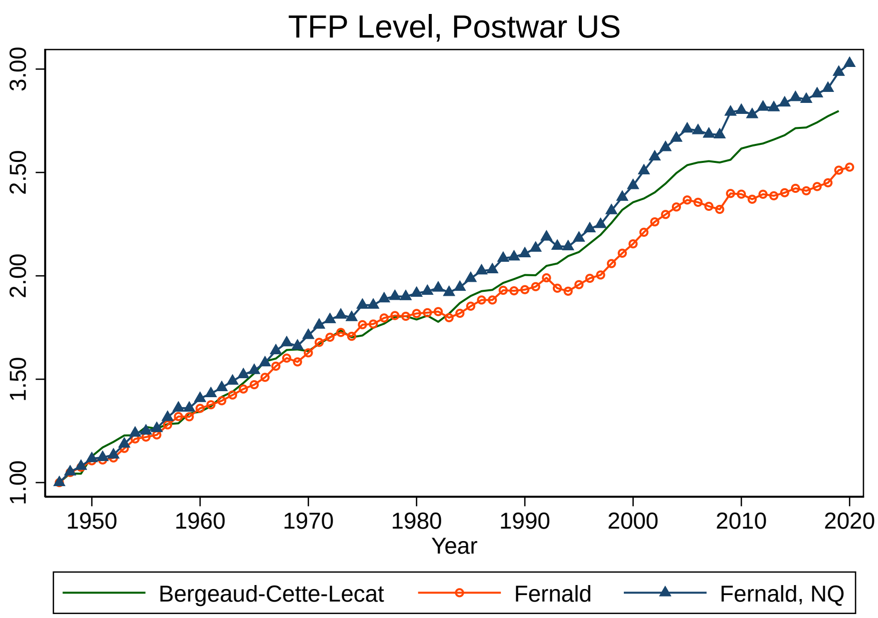

Notes: TFP levels, A^{bcl}, A_t^f, and A_t^n. Data from Fernald (2012) and Bergeaud et al. (2016).

B Why In-Sample Tests Are Meaningless

Figure 13 shows why in-sample fit cannot be used to test the models' predictions. The figure uses the most basic measure, GDP per capita, which is often used to argue for exponential growth. The left panel is a plot used in textbooks to argue that GDP growth is exponential but in fact the figure contains little information. The right panel compares the in-sample fit of the exponential curve with that of a quadratic curve. Visually both appears to fit reasonably well and if anything the quadratic fit is better (1% higher R^2). The main point however, is that we learn nothing from this exercise beyond the fact that GDP per capita is convex, which is true of any linear TFP model. Another point, is that the exponential model makes increasingly larger errors at the end. The log-scale on the left makes it look like the mistakes are small, but in fact the fit is off by about \$8k at the end of the sample, which is 13% of the actual value (around \$60k).

Figure 13: US GDP per Capita (2010 \$ppp)

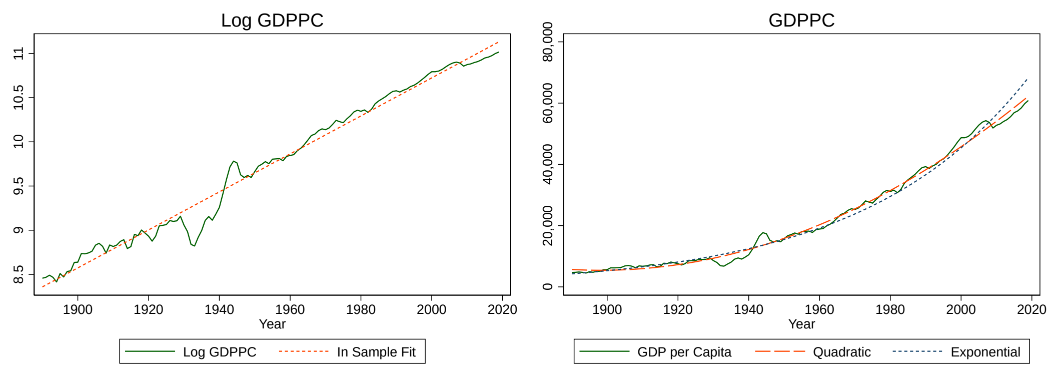

Notes: Bergeaud et al. (2016).