Tracking the Slowdown in Long-Run GDP Growth*

Juan Antolin-Diaz — Fulcrum Asset Management

Thomas Drechsel — LSE, CFM

Ivan Petrella — Bank of England, Birkbeck, CEPR

This draft: April 20, 2016   First draft: October 24, 2014

Abstract

Using a dynamic factor model that allows for changes in both the long-run growth rate of output and the volatility of business cycles, we document a significant decline in long-run output growth in the United States. Our evidence supports the view that most of this slowdown occurred prior to the Great Recession. We show how to use the model to decompose changes in long-run growth into its underlying drivers. At low frequencies, a decline in the growth rate of labor productivity appears to be behind the recent slowdown in GDP growth for both the US and other advanced economies. When applied to real-time data, the proposed model is capable of detecting shifts in long-run growth in a timely and reliable manner.

Keywords: Long-run growth; Business cycles; Productivity; Dynamic factor models; Real-time data.

JEL Classification Numbers: E32, E23, O47, C32, E01.

*Antolin-Diaz: Department of Macroeconomic Research, Fulcrum Asset Management, London W1H 5BT, UK; E-Mail: juan.antolin-diaz@fulcrumasset.com. Drechsel: Department of Economics and Centre for Macroeconomics, LSE, London, WC2A 2AE, UK; E-Mail: t.a.drechsel@lse.ac.uk. Petrella: Structural Economic Analysis Division, Bank of England, London EC2R 8AH, UK; E-Mail: Ivan.Petrella@bankofengland.co.uk. The views expressed in this paper are solely the responsibility of the authors and cannot be taken to represent those of Fulcrum Asset Management or the Bank of England, or to state Bank of England policy. This paper should not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee or Financial Policy Committee. We thank Neele Balke, James Bell, Gavyn Davies, Davide Delle Monache, Wouter den Haan, Robert Gordon, Andrew Harvey, Ed Hoyle, Dimitris Korobilis, Juan Rubio-Ramirez, Ron Smith, James Stock, Paolo Surico, Jon Temple, Silvana Tenreyro, Joaquin Vespignani and the seminar participants at Banca d'Italia, Birkbeck Centre for Applied Macro, ECARES, European Central Bank, BBVA Research, the Applied Econometrics Workshop at the Federal Reserve Bank of St. Louis, the World Congress of the Econometric Society, Norges Bank and LSE for useful comments and suggestions. Alberto D'Onofrio provided excellent research assistance.

1 Introduction

"The global recovery has been disappointing (...) Year after year we have had to explain from mid-year on why the global growth rate has been lower than predicted as little as two quarters back". Stanley Fischer, August 2014.

The slow pace of the recovery from the Great Recession of 2007-2009 has prompted questions about whether the long-run growth rate of GDP in advanced economies is lower now than it has been on average over the past decades (see e.g. Fernald, 2014, Gordon, 2014b, Summers, 2014). Indeed, forecasts of US and global real GDP growth have been persistently too optimistic for the last six years.[1] As emphasized by Orphanides (2003), real-time misperceptions about the long-run growth of the economy can play a large role in monetary policy mistakes. Moreover, small changes in assumptions about the long-run growth rate of output can have large implications on fiscal sustainability calculations (Auerbach, 2011). This calls for a framework that takes the uncertainty about long-run growth seriously and can inform decision-making in real time. In this paper, we present a dynamic factor model (DFM) which allows for gradual changes in the mean and the variance of real output growth. By incorporating a broad panel of economic activity indicators, DFMs are capable of precisely estimating the cyclical comovement in macroeconomic data in a real-time setting. Our model exploits this to track changes in the long-run growth rate of real GDP in a timely and reliable manner, separating them from their cyclical counterpart.[2]

The evidence of a decline in long-run US growth is accumulating, as documented

[1]For instance, Federal Open Market Committee (FOMC) projections since 2009 expected US growth to accelerate substantially, only to downgrade the forecast back to 2% throughout the course of the subsequent year. An analysis of forecasts produced by international organizations and private sector economists reveals the same pattern, see Pain et al. (2014) for a retrospective.

[2]Throughout this paper, our concept of the long run refers to changes in growth that are permanent in nature, i.e. do not mean-revert, as in Beveridge and Nelson (1981). In practice this should be thought of as frequencies lower than the business cycle.

by the recent growth literature such as Fernald and Jones (2014). Lawrence Summers and Robert Gordon have articulated a particularly pessimistic view of long-run growth which contrasts with the optimism prevailing before the Great Recession (see Jorgenson et al., 2006). To complement this evidence, we start our analysis by presenting the results of two popular structural break tests proposed by Nyblom (1989) and Bai and Perron (1998). Both suggest that a possible shift in the mean of US real GDP growth exists, the latter approach suggesting that a break probably occurred in the early part of the 2000's.[3] However, sequential testing using real-time data reveals that the break would not have been detected at conventional significance levels until as late as mid-2012, highlighting the problems of conventional break tests for real-time analysis (see also Benati, 2007). To address this issue, we introduce two novel features into an otherwise standard DFM of real activity data. First, we allow the mean of real GDP growth, and possibly other series, to drift gradually over time. As emphasized by Cogley (2005), if the long-run output growth rate is not constant, it is optimal to give more weight to recent data when estimating its current state. By taking a Bayesian approach, we can combine our prior beliefs about the rate at which the past information should be discounted with the information contained in the data. We also characterize the uncertainty around estimates of long-run growth taking into account both filtering and parameter uncertainty. Second, we allow for stochastic volatility (SV) in the innovations to both factors and idiosyncratic components. Given our interest in studying the entire postwar period, the inclusion of SV is essential to capture the substantial changes in the volatility of output that have taken place in this sample, such as the "Great Moderation" first reported by Kim and Nelson (1999a) and McConnell and Perez-Quiros (2000), as well as the cyclicality of macroeconomic volatility as documented by Jurado et al. (2014).

[3] This finding is consistent with the analysis of US real GDP by Luo and Startz (2014), as well as Fernald (2014), who applies the Bai and Perron (1998) test to US labor productivity.

When applied to US data, our model concludes that long-run GDP growth declined meaningfully during the 2000's and currently stands at about 2%, more than one percentage point lower than the postwar average. The results are supportive of a gradual decline rather than a discrete break. Since in-sample results obtained with revised data often underestimate the uncertainty faced by policymakers in real time, we repeat the exercise using real-time vintages of data. The model detects the fall from the beginning of the 2000's onwards, and by the summer of 2010 it reaches the significant conclusion that a decline in long-run growth is behind the slow recovery, well before the structural break tests become conclusive.

We also investigate the performance of the model in "nowcasting" short-term developments in GDP. Since the seminal contributions of Evans (2005) and Giannone et al. (2008) DFMs have become the standard tool for this purpose.[4] Interestingly, our analysis shows that standard DFM forecasts revert very quickly to the unconditional mean of GDP, so taking into account the variation in long-run GDP growth substantially improves point and density GDP forecasts even at very short horizons.

Finally, we extend our model in order to disentangle the drivers of secular fluctuations of GDP growth. Edge et al. (2007) emphasize the relevance as well as the difficulty of tracking permanent shifts in productivity growth in real time. In our framework, long-run output growth can be decomposed into labor productivity and labor input trends. The results of this decomposition exercise point to a slowdown in labor productivity as the main driver of recent weakness in GDP growth. Applying the model to other advanced economies, we provide evidence that the weakening in labor productivity appears to be a global phenomenon.

Our work is closely related to two strands of literature. The first one encompasses papers that allow for structural changes within the DFM framework. Del Negro and

[4] An extensive survey of the nowcasting literature is provided by Banbura et al. (2012), who also demonstrate, in a real-time context, the good out-of-sample performance of DFM nowcasts.

Otrok (2008) model time variation in factor loadings and volatilities, while Marcellino et al. (2014) show that the addition of SV improves the performance of the model for short-term forecasting of euro area GDP.[5] Acknowledging the importance of allowing for time-variation in the means of the variables, Stock and Watson (2012) pre-filter their data set in order to remove any low-frequency trends from the resulting growth rates using a biweight local mean. In his comment to their paper, Sims (2012) suggests to explicitly model, rather than filter out, these long-run trends, and emphasizes the importance of evolving volatilities for describing and understanding macroeconomic data. We see the present paper as extending the DFM literature, and in particular its application to tracking GDP, in the direction suggested by Chris Sims. The second strand of related literature takes a similar approach to decomposing long-run GDP growth into its drivers, in particular Gordon (2010, 2014a) and Reifschneider et al. (2013). Relative to these studies, we emphasize the importance of using a broader information set, as well as a Bayesian approach, which allows to use priors to inform the estimate of long-run growth, and to characterize the uncertainty around the estimate stemming both from filtering and parameter uncertainty.

The remainder of this paper is organized as follows. Section 2 presents preliminary evidence of a slowdown in long-run US GDP growth. Section 3 discusses the implications of time-varying long-run output growth and volatility for DFMs and presents our model. Section 4 applies the model to US data and documents the decline in long-run growth. The implications for tracking GDP in real time as well as the key advantages of our methodology are discussed. Section 5 decomposes the changes in long-run output growth into its underlying drivers. Section 6 concludes.

[5] While the model of Del Negro and Otrok (2008) includes time-varying factor loadings, the means of the observable variables are still treated as constant.

2 Preliminary Evidence

The literature on economic growth favors a view of the long-run growth rate as a process that evolves over time. It is by now widely accepted that a slowdown in productivity and long-run output growth occurred in the early 1970's, and that accelerating productivity in the IT sector led to a boom in the late 1990's.[6] In contrast, in the context of econometric modeling the possibility that long-run growth is time-varying is the source of a long-standing controversy. In their seminal contribution, Nelson and Plosser (1982) model the (log) level of real GDP as a random walk with drift. This implies that after first-differencing, the resulting growth rate fluctuates around a constant mean, an assumption still embedded in many econometric models. After the slowdown in productivity became apparent in the 1970's, many researchers such as Clark (1987) modeled the drift term as an additional random walk, implying that the level of GDP is integrated of order two. The latter assumption would also be consistent with the local linear trend model of Harvey (1985), the Hodrick and Prescott (1997) filter, and Stock and Watson (2012)'s practice of removing a local biweight mean from the growth rates before estimating a DFM. The $I(2)$ assumption is nevertheless controversial since it implies that the growth rate of output can drift without bound. Consequently, papers such as Perron and Wada (2009), have modeled the growth rate of GDP as stationary around a trend with one large break around 1973.

Ever since the Great Recession of 2007-2009 US real GDP has grown well below its postwar average, once again raising the question whether its mean may have declined. There are two popular strategies that could be followed from a frequentist perspective to detect parameter instability or the presence of breaks in the mean growth rate. The first one is Nyblom's (1989) L-test as described in Hansen (1992), which tests

[6] For a retrospective on the productivity slowdown, see Nordhaus (2004). Oliner and Sichel (2000) provide evidence on the role of the IT sector in the acceleration of the late 1990's.

Figure 1: Real-Time Test Statistics of the Nyblom and Bai-Perron Tests

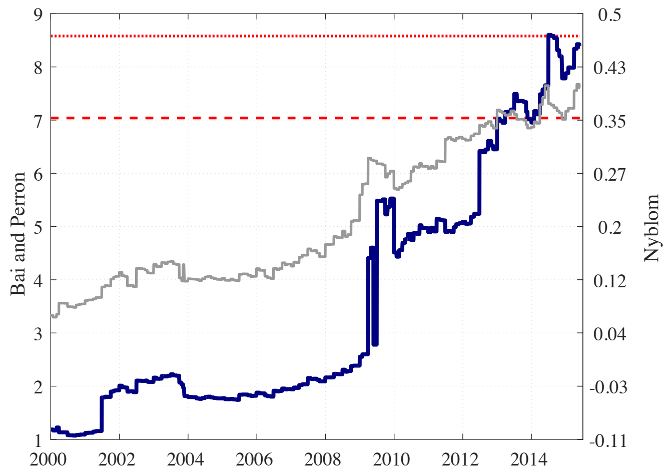

Note: The gray and blue lines are the values of the test statistics obtained from sequentially re-applying the Nyblom (1989) and Bai and Perron (1998) tests in real time as new National Accounts vintages are being published. In both cases, the sample starts in 1947:Q2 and the test is re-applied for every new data release occurring after the beginning of 2000. The dotted and dashed red lines represent the 5% and 10% critical values corresponding to the two tests.

the null hypothesis of constant parameters against the alternative that the parameters follow a martingale. Modeling real GDP growth as an AR(1) over the sample 1947-2015 this test rejects the stability of the constant term at the 10% significance level.[7] The second commonly used approach, which can determine the number and timing of multiple discrete breaks, is the Bai and Perron (1998) test. This test finds evidence in favor of a single break in the mean of US real GDP growth at the 10%-level. The most likely break date is in the second quarter of 2000. In related research, Fernald (2014) provides evidence for breaks in labor productivity in 1973:Q2, 1995:Q3, and 2003:Q1, and links the latter two to developments in the IT sector. From a Bayesian perspective, Luo and Startz (2014) calculate the posterior probability of a single break and find the most likely break date to be 2006:Q1 for the full postwar sample and 1973:Q1 for a

[7]The same result holds for an AR(2) specification. In both cases, stability of the autoregressive coefficients cannot be rejected, whereas stability of the variance is rejected at the 1%-level. Appendix B provides the full results of both tests applied in this section. The appendix to the paper is available at: `http://personal.lse.ac.uk/drechsel/papers/ADP_appendix.pdf`.

sample excluding the 2000's.

The above results indicate that substantial evidence for a recent change in the mean of US GDP growth has built up. However, the strategy of applying conventional tests and introducing deterministic breaks into econometric models is not satisfactory for the purposes of real-time decision making. In fact, the detection of change in the mean of GDP growth can arrive with substantial delay. To demonstrate this, a sequential application of the Nyblom (1989) and Bai and Perron (1998) tests using real-time data is presented in Figure 1. The evolution of the test statistics in real-time reveals that a break would not have been detected at the 10% significance levels until as late as mid-2012, which is more than ten years later than the actual break date suggested by the Bai and Perron (1998) procedure. The Nyblom (1989) test, which is designed to detect gradual change rather than a discrete break, becomes significant roughly at the same time. This lack of timeliness highlights the importance of an econometric framework capable of quickly adapting to changes in long-run growth as new information arrives.

3 Econometric Framework

DFMs in the spirit of Geweke (1977), Stock and Watson (2002) and Forni et al. (2009) capture the idea that a small number of unobserved factors drives the comovement of a possibly large number of macroeconomic time series, each of which may be contaminated by measurement error or other sources of idiosyncratic variation. Their theoretical appeal (see e.g. Sargent and Sims, 1977 or Giannone et al., 2006), as well as their ability to parsimoniously model large data sets, have made them a workhorse of empirical macroeconomics. Giannone et al. (2008) and Banbura et al. (2012) have pioneered the use of DFMs to produce current-quarter forecasts ("nowcasts") of GDP growth by exploiting more timely monthly indicators and the factor structure of the

data. Given the widespread use of DFMs to track GDP in real time, this paper aims to make these models robust to changes in long-run growth. We do so by introducing two novel features into the DFM framework. First, we allow the long-run growth rate of real GDP, and possibly other series, to vary over time. Second, we allow for stochastic volatility (SV) in the innovations to both factors and idiosyncratic components, given our interest in studying the entire postwar period for which drastic changes in volatility have been documented. With these changes, the DFM proves to be a powerful tool to detect changes in long-run growth. The information contained in a broad panel of activity indicators facilitates the timely decomposition of real GDP growth into persistent long-run movements, cyclical fluctuations and short-lived noise.

3.1 The Model

Let $\mathbf{y}_t$ be an $n \times 1$ vector of observable macroeconomic time series, and let $\mathbf{f}_t$ denote a $k \times 1$ vector of latent common factors. It is assumed that $n >> k$, i.e. the number of observables is much larger than the number of factors. Formally,

$$\mathbf{y}_t = \mathbf{c}_t + \mathbf{\Lambda} \mathbf{f}_t + \mathbf{u}_t, \tag{1}$$

where $\mathbf{\Lambda}$ contains the loadings on the common factors and $\mathbf{u}_t$ is a vector of idiosyncratic components.[8] Shifts in the long-run mean of $\mathbf{y}_t$ are captured by time-variation in $\mathbf{c}_t$. In principle one could allow time-varying intercepts in all or a subset of the variables in the system. Moreover, time variation in a given series could be shared by other series.

[8]The model can be easily extended to include lags of the factor in the measurement equation. In the latter case, it is sensible to avoid overfitting by choosing priors that shrink the additional lag coefficients towards zero (see D'Agostino et al., 2015, and Luciani and Ricci, 2014). We consider this possibility when we explore robustness of our results to using larger data panels in Section 4.6.

$\mathbf{c}_t$ is therefore flexibly specified as

$$\mathbf{c}_t = \begin{bmatrix} \mathbf{B} & \mathbf{0} \\ \mathbf{0} & \mathbf{c} \end{bmatrix} \begin{bmatrix} \mathbf{a}_t \\ 1 \end{bmatrix}, \tag{2}$$

where $\mathbf{a}_t$ is an $r \times 1$ vector of time-varying means, $\mathbf{B}$ is an $m \times r$ matrix which governs how the time-variation affects the corresponding observables, and $\mathbf{c}$ is an $(n-m) \times 1$ vector of constants. In our baseline specification, $\mathbf{a}_t$ will be a scalar capturing time-variation in long-run real GDP growth, which is shared by real consumption growth, so that $r = 1, m = 2$. A detailed discussion of this and additional specifications of $\mathbf{c}_t$ will be provided in Section 3.2.

Throughout the paper, we focus on the case of a single dynamic factor by setting $k = 1$ (i.e. $\mathbf{f}_t = f_t$).[9] The laws of motion of the latent factor and the idiosyncratic components are

$$(1 - \phi(L))f_t = \sigma_{\varepsilon_t}\varepsilon_t, \tag{3}$$

$$(1 - \rho_i(L))u_{i,t} = \sigma_{\eta_{i,t}}\eta_{i,t}, \quad i = 1, \ldots, n \tag{4}$$

where $\phi(L)$ and $\rho_i(L)$ denote polynomials in the lag operator of order p and q, respectively. The idiosyncratic components are cross-sectionally orthogonal and are assumed to be uncorrelated with the common factor at all leads and lags, i.e. $\varepsilon_t \overset{iid}{\sim} N(0,1)$ and $\eta_{i,t} \overset{iid}{\sim} N(0,1)$.

Finally, the dynamics of the model's time-varying parameters are specified to follow

[9]For the purpose of tracking real GDP with a large number of closely related activity indicators, the use of one factor is appropriate, which is explained in more detail in Sections 4.1 and 4.2. Also note that we order real GDP growth as the first element of $\mathbf{y}_t$, and normalize the loading for GDP to unity. This serves as an identifying restriction in our estimation algorithm. Bai and Wang (2015) discuss minimal identifying assumptions for DFMs.

driftless random walks:

$$
\begin{aligned}
a_{j,t} &= a_{j,t-1} + v_{a_{j,t}}, & v_{a_{j,t}} &\overset{iid}{\sim} N(0, \omega^2_{a,j}) \quad j = 1, \dots, r & (5)\\
\log \sigma_{\varepsilon_t} &= \log \sigma_{\varepsilon_{t-1}} + v_{\varepsilon,t}, & v_{\varepsilon,t} &\overset{iid}{\sim} N(0, \omega^2_{\varepsilon}) & (6)\\
\log \sigma_{\eta_{i,t}} &= \log \sigma_{\eta_{i,t-1}} + v_{\eta_{i,t}}, & v_{\eta_{i,t}} &\overset{iid}{\sim} N(0, \omega^2_{\eta,i}) \quad i = 1, \dots, n & (7)
\end{aligned}
$$

where $a_{j,t}$ are the r time-varying elements in $\mathbf{a}_t$, and σ_{ε_t} and $\sigma_{\eta_{i,t}}$ capture the SV of the innovations to factor and idiosyncratic components. Our motivation for specifying the time-varying parameters as random walks is similar to Primiceri (2005). While in principle it is unrealistic model real GDP growth as a process that could wander in an unbounded way, as long as the variance of the process is small and the drift is considered to be operating for a finite period of time, the assumption is innocuous. Moreover, modeling a trend as a random walk is more robust to misspecification when the actual process is instead characterized by discrete breaks, whereas models with discrete breaks might not be robust to the true process being a random walk.[10] Finally, the random walk assumption also has the desirable feature that, unlike stationary models, confidence bands around forecasts of real GDP growth increase with the forecast horizon, reflecting uncertainty about the possibility of future breaks or drifts in long-run growth.

Note that a standard DFM is usually specified under two assumptions. First, the original data have been differenced appropriately so that both the factor and the idiosyncratic components can be assumed to be stationary. Second, it is assumed that the innovations in the idiosyncratic and common components are iid. In equations (1)-(7) we have relaxed these assumptions to allow for two novel features, a stochastic

[10]We demonstrate this point with the use of Monte Carlo simulations, showing that a random walk trend in real GDP growth ‘learns’ quickly about a discrete break once it has occurred. On the other hand, the random walk does not detect a drift when there is not one, despite the presence of a large cyclical component. Appendix C provides a discussion and the full results of these simulations.

trend in the mean of selected series, and SV. By shutting down these features, we can recover the specifications previously proposed in the literature, which are nested in our framework. We obtain the DFM with SV of Marcellino et al. (2014) if we shut down time-variation in the intercepts of the observables, i.e. set $r = m = 0$ and $\mathbf{c}_t = \mathbf{c}$. If we further shut down the SV, i.e. set $\omega_{a,j}^2 = \omega_\epsilon^2 = \omega_{\eta,i}^2 = 0$, we obtain the specification of Banbura and Modugno (2014) and Banbura et al. (2012).

3.2 A Baseline Specification for Long-Run Growth

Equations (1) and (2) allow for stochastic trends in the mean of all or a subset of selected observables in $\mathbf{y}_t$. This paper focuses on tracking changes in the long-run growth rate of real GDP. For this purpose, the simplest specification of $\mathbf{c}_t$ is to include a time-varying intercept only in GDP and to set $\mathbf{B} = 1$. However, a number of empirical studies (e.g. Harvey and Stock, 1988, Cochrane, 1994, and Cogley, 2005) argue that incorporating information about consumption is informative about the permanent component in GDP as predicted by the permanent income hypothesis. The theory predicts that consumers, smoothing consumption throughout their lifetime, should react more strongly to permanent, as opposed to transitory, changes in income. As a consequence, looking at GDP and consumption data together will help separating growth into long-run and cyclical fluctuations.[11] Therefore, our baseline specification imposes that consumption and output grow at the same rate g_t in the long-run. On the contrary, we do not impose that investment also grows at this rate, as would be the case in the basic neoclassical growth model, since the presence of investment-specific technological change implies that real investment has a different low-frequency trend (Greenwood et al., 1997).

[11] While a strict interpretation of the permanent income hypothesis is rejected in the data, from an econometric point of view the statement applies as long as permanent changes are the main driver of consumption. See Cochrane (1994) for a very similar discussion.

Formally, ordering real GDP and consumption growth first, and setting $m = 2$ and $r = 1$, this is represented as

$$\mathbf{a}_t = g_t, \quad \mathbf{B} = [1\ 1]' \tag{8}$$

Note that in this baseline specification we model time-variation only in the intercept for GDP and consumption while leaving it constant for the other observables. Of course it may be the case that some of the remaining $n - m$ series in $\mathbf{y}_t$ feature low frequency variation in their means. For instance, as mentioned above, this could be the case for investment. The key question is whether leaving it unspecified will affect the estimate of the long-run growth rate of GDP, which is our main object of interest. We ensure that this is not the case by allowing for persistence (and, in particular, we do not rule out unit roots) in the idiosyncratic components. If a series does feature a unit root which is not included in $\mathbf{a}_t$, its trend component will be absorbed by the idiosyncratic component. The choice of which elements to include in $\mathbf{a}_t$ therefore reflects the focus of a particular application.[12] Of course, if two series share the same underlying low-frequency component, and this is known with certainty, explicitly accounting for the shared low frequency variation will improve the precision of the estimation, but the risk of incorrectly including the trend is much larger than the risk of incorrectly excluding it. Therefore, in our baseline specification we include in $\mathbf{a}_t$ the intercept for GDP and consumption, while leaving any possible low-frequency variation in other series to be captured by the respective idiosyncratic components.[13]

[12] In principle, these unmodeled trends could still be recovered from our specification by applying a Beveridge-Nelson decomposition to its estimated idiosyncratic component. In practice, any low-frequency variation in the idiosyncratic component is likely to be obscured by a large amount of high frequency noise in the data and as result the extracted Beveridge-Nelson trend component will be imprecisely estimated, and as Morley et al. (2003) show, will not be smooth. In our specification, the elements of $\mathbf{a}_t$ are instead extracted directly, so that we are able to improve the extraction by imposing additional assumptions (e.g. smoothness) and prior beliefs (e.g. low variability) on its properties.

[13] We confirm this line of reasoning with a series of Monte Carlo experiments, in which data is generated from a system that features low-frequency movements in more series, which are left unmodeled in the estimation. Both in the case of series with independent trends and the case of series which share the trend of interest, the fact that they are left unmodeled has little impact on the estimate of

An extension to include additional time-varying intercepts is straightforward through the flexible construction of $\mathbf{c}_t$ in equation (2). In fact, in Section 5 we explore how interest in the low-frequency movements of additional series leads to alternative choices for $\mathbf{a}_t$ and $\mathbf{B}$.[14]

3.3 Dealing with Mixed Frequencies and Missing Data

Tracking activity in real time requires a model that can efficiently incorporate information from series measured at different frequencies. In particular, it must include both quarterly variables, such as the growth rate of real GDP, as well as more timely monthly indicators of real activity. Therefore, the model is specified at monthly frequency, and following Mariano and Murasawa (2003), the (observed) quarterly growth rates of a generic quarterly variable, x_t^q, can be related to the (unobserved) monthly growth rate x_t^m and its lags using a weighted mean. Specifically,

$$x_t^q = \frac{1}{3}x_t^m + \frac{2}{3}x_{t-1}^m + x_{t-2}^m + \frac{2}{3}x_{t-3}^m + \frac{1}{3}x_{t-4}^m, \tag{9}$$

and only every third observation of x_t^q is actually observed. Substituting the corresponding line of (1) into (9) yields a representation in which the quarterly variable depends on the factor and its lags. The presence of mixed frequencies is thus reduced to a problem of missing data in a monthly model.

Besides mixed frequencies, additional sources of missing data in the panel include: the "ragged edge" at the end of the sample, which stems from the non-synchronicity of data releases; missing data at the beginning of the sample, since some data series

the latter. Appendix C presents further discussion and the full results of these simulations.

[14] Note that the limiting case explicitly models time-varying intercept in all indicators, so that $m = r = n$ and $\mathbf{B} = \mathbf{I}_n$, i.e. an identity matrix of dimension n. See Creal et al. (2010) and Fleischman and Roberts (2011) for similar approaches. This setup would imply that the number of state variables increases with the number of observables, which severely increases the computational burden of the estimation, while offering little additional evidence with respect to the focus of this paper.

have been created or collected more recently than others; and missing observations due to outliers and data collection errors. Our Bayesian estimation method exploits the state space representation of the DFM and jointly estimates the latent factors, the parameters, and the missing data points using the Kalman filter (see Durbin and Koopman, 2012, for a textbook treatment).

3.4 State Space Representation and Estimation

The model features autocorrelated idiosyncratic components (see equation (4)). In order to cast it in state-space form, we include the idiosyncratic components of the quarterly variables in the state vector, and we redefine the system for the monthly indicators in terms of quasi-differences (see e.g. Kim and Nelson, 1999b, pp. 198-199, and Bai and Wang, 2015).[15] The model is estimated with Bayesian methods simulating the posterior distribution of parameters and factors using a Markov Chain Monte Carlo (MCMC) algorithm. We closely follow the Gibbs-sampling algorithm for DFMs proposed by Bai and Wang (2015), but extend it to include mixed frequencies, the time-varying intercept, and SV. The SVs are sampled using the approximation of Kim et al. (1998), which is considerably faster than the exact Metropolis-Hastings algorithm of Jacquier et al. (2002). Our complete sampling algorithm together with the details of the state space representation can be found in Appendix D.

[15] Since the quarterly variables are observed only every third month, we cannot take the quasi-difference for their idiosyncratic components, which are instead added as an additional state with the corresponding transition dynamics. Banbura and Modugno (2014) suggest including all of the idiosyncratic components as additional elements of the state vector. Our solution has the desirable feature that the number of state variables will increase with the number of quarterly variables, rather than the total number of variables, leading to a gain of computational efficiency.

4 Results for US Data

4.1 Data Selection

Our data set includes four key business cycle variables measured at quarterly frequency (output, consumption, investment and aggregate hours worked), as well as a set of 24 monthly indicators which are intended to provide additional information about cyclical developments in a timely manner.

The included quarterly variables are strongly procyclical and are considered key indicators of the business cycle (see e.g. Stock and Watson, 1999). Furthermore, theory predicts that they will be useful in disentangling low frequency movements from cyclical fluctuations in output growth. Indeed, as discussed in Section 3.2, the permanent income hypothesis predicts that consumption data will be particularly useful for the estimation of the long-run growth component, g_t.[16] On the other hand, investment and hours worked are very sensitive to cyclical fluctuations, and thus will be particularly informative for the estimation of the common factor, f_t.[17]

The additional monthly indicators are crucial to our objective of disentangling in real time the cyclical and long-run components of GDP growth, since the quarterly variables are only available with substantial delay. In principle, a large number of candidate series are available to inform the estimate of f_t, and indirectly, of g_t. In practice,

[16] Due to the presence of faster technological change in the durable goods sector there is a downward trend in the relative price of durable goods. As a consequence, measured consumption grows faster than overall GDP. Following a long tradition in the literature (see e.g. Whelan, 2003), we construct a Fisher index of non-durables and services and use its growth rate as an observable variable in the panel. It can be verified that the ratio of consumption defined in this manner to real GDP displays no trend in the data, unlike the trend observed in the ratio of overall consumption to GDP.

[17] We define investment as a chain-linked aggregate of business fixed investment and consumption of durable goods, which is consistent with our treatment of consumption. In order to obtain a measure of hours for the total economy, we follow the methodology of Ohanian and Raffo (2012) and benchmark the quarterly series of hours in the non-farm business sector provided by the BLS to the annual estimates of hours in the total economy compiled by the Conference Boards Total Economy Database (TED). The TED series has the advantage of being comparable across countries (Ohanian and Raffo, 2012), which will be useful for our international results in Section 5.

however, macroeconomic data series are typically clustered in a small number of broad categories (such as production, employment, or income) for which disaggregated series are available along various dimensions (such as economic sectors, demographic characteristics, or expenditure categories). The choice of which available series to include for estimation can therefore be broken into, first, a choice of which broad categories to include, and second, to which level and along which dimensions of disaggregation.

With regards to which broad categories of data to include, previous studies agree that prices, monetary and financial indicators are uninformative for the purpose of tracking real GDP, and argue for extracting a single common factor that captures real economic activity.[18] As for the possible inclusion of disaggregated series within each category, Boivin and Ng (2006) argue that the presence of strong correlation in the idiosyncratic components of disaggregated series of the same category will be a source of misspecification that can worsen the performance of the model in terms of in-sample fit and out-of-sample forecasting of key series.[19] Alvarez et al. (2012) investigate the trade-off between DFMs with very few indicators, where the good large-sample properties of factor models are unlikely to hold, and those with a very large amount of indicators, where the problems above are likely to arise. They conclude that using a medium-sized panel with representative indicators of each category yields the best forecasting results.

The above considerations lead us to select 24 monthly indicators that include the high-level aggregates for all of the available broad categories that capture real activity, without overweighting any particular category. The complete list of variables contained in our data set is presented in Table 1. As the table shows, we include representative series of expenditure and income, the labor market, production and sales, foreign trade,

[18]Giannone et al. (2005) conclude that that prices and monetary indicators do not contribute to the precision of GDP nowcasts. Banbura et al. (2012), Forni et al. (2003) and Stock and Watson (2003) find at best mixed results for financial variables.

[19]This problem is exacerbated by the fact that more detailed disaggregation levels and dimensions are available for certain categories of data, such as employment, meaning that the disaggregation will automatically 'tilt' the factor estimates towards that category.

housing and business and consumer confidence.[20] The inclusion of all the available monthly surveys is particularly important. Apart from being the most timely series available, these are unlikely to feature permanent shifts in their mean by construction, and have a high signal-to-noise ratio. They thus provide a clean signal to separate the cyclical component of GDP growth from its long-run counterpart. In Section 4.6 we explore sensitivity of our results to the size and composition of the data panel used.

Our panel spans the period January 1947 to March 2015. The start of our sample coincides with the year for which quarterly national accounts data are available from the Bureau of Economic Analysis. This enables us to study the evolution of long-run growth over the entire postwar period.[21]

4.2 Model Settings and Priors

The choice of the data set justifies the single-factor structure of the model. f_t can in this case be interpreted as a coincident indicator of real economic activity (see e.g. Stock and Watson, 1989, and Mariano and Murasawa, 2003). The number of lags in the polynomials $\phi(L)$ and $\rho(L)$ is set to $p = 2$ and $q = 2$ as in Stock and Watson (1989). We wish to impose as little prior information as possible, so we use uninformative priors for the factor loadings and the autoregressive coefficients of factors and idiosyncratic components. The variances of the innovations to the time-varying parameters, namely ω_a^2, ω_ε^2 and $\omega_{\eta,i}^2$ in equations (5)-(7) are however difficult to identify from the information contained in the likelihood alone. As the literature on Bayesian

[20]When there are multiple candidates for the high-level aggregate of a category, we include both. For example, we include employment as measured both by the establishment and household surveys, and consumer confidence as surveyed both by the Conference Board and the University of Michigan.

[21]We take full advantage of the Kalman filter's ability to deal with missing observations at any point in the sample, and we are able to incorporate series that become available substantially later than 1947, up to as late as 2007. Note that for consumption expenditures, monthly data became available in 1959, whereas quarterly data is available from 1947. In order to use all available data, we apply the polynomial in Equation (9) to the monthly data and treat the series as quarterly, with available observations for the last month of the quarter for 1947-1958 and for all months since 1959.

Table 1: Data series used in empirical analysis

	Type	Start Date	Transformation	Publ. Lag
Quarterly time series				
Real GDP	Expenditure & Income	Q2:1947	% QoQ Ann.	26
Real Consumption (excl. durables)	Expenditure & Income	Q2:1947	% QoQ Ann.	26
Real Investment (incl. durable cons.)	Expenditure & Income	Q2:1947	% QoQ Ann.	26
Total Hours Worked	Labor Market	Q2:1948	% QoQ Ann.	28
Monthly indicators				
Real Personal Income less Trans. Paym.	Expenditure & Income	Feb 59	% MoM	27
Industrial Production	Production & Sales	Jan 47	% MoM	15
New Orders of Capital Goods	Production & Sales	Mar 68	% MoM	25
Real Retail Sales & Food Services	Production & Sales	Feb 47	% MoM	15
Light Weight Vehicle Sales	Production & Sales	Feb 67	% MoM	1
Real Exports of Goods	Foreign Trade	Feb 68	% MoM	35
Real Imports of Goods	Foreign Trade	Feb 69	% MoM	35
Building Permits	Housing	Feb 60	% MoM	19
Housing Starts	Housing	Feb 59	% MoM	26
New Home Sales	Housing	Feb 63	% MoM	26
Payroll Empl. (Establishment Survey)	Labor Market	Jan 47	% MoM	5
Civilian Empl. (Household Survey)	Labor Market	Feb 48	% MoM	5
Unemployed	Labor Market	Feb 48	% MoM	5
Initial Claims for Unempl. Insurance	Labor Market	Feb 48	% MoM	4
Monthly indicators (soft)				
Markit Manufacturing PMI	Business Confidence	May 07	-	-7
ISM Manufacturing PMI	Business Confidence	Jan 48	-	1
ISM Non-manufacturing PMI	Business Confidence	Jul 97	-	3
NFIB: Small Business Optimism Index	Business Confidence	Oct 75	Diff 12 M.	15
U. of Michigan: Consumer Sentiment	Consumer Confidence	May 60	Diff 12 M.	-15
Conf. Board: Consumer Confidence	Consumer Confidence	Feb 68	Diff 12 M.	-5
Empire State Manufacturing Survey	Business (Regional)	Jul 01	-	-15
Richmond Fed Mfg Survey	Business (Regional)	Nov 93	-	-5
Chicago PMI	Business (Regional)	Feb 67	-	0
Philadelphia Fed Business Outlook	Business (Regional)	May 68	-	0

Notes: % QoQ Ann. refers to the quarter on quarter annualized growth rate, % MoM refers to $(y_t - y_{t-1})/y_{t-1}$ while Diff 12 M. refers to $y_t - y_{t-12}$. The last column shows the average publication lag, i.e. the number of days elapsed from the end of the period that the data point refers to until its publication by the statistical agency. All series were obtained from the Haver Analytics database.

VARs documents, attempts to use non-informative priors for these parameters will in many cases produce posterior estimates which imply a relatively large amount of time-variation. While this will tend to improve the in-sample fit of the model it is also likely to worsen out-of-sample forecast performance. We therefore use priors to shrink these variances towards zero, i.e. towards the standard DFM which excludes time-varying long-run GDP growth and SV. In particular, for ω_a^2 we set an inverse gamma prior with one degree of freedom and scale equal to 0.001.[22] For ω_ϵ^2 and $\omega_{\eta,i}^2$ we set an inverse

[22]To gain an intuition about this prior, note that over a period of ten years, this would imply that the random walk process of the long-run growth rate is expected to vary with a standard deviation

gamma prior with one degree of freedom and scale equal to 0.0001, closely following Cogley and Sargent (2005) and Primiceri (2005).[23] We estimate the model with 7000 replications of the Gibbs-sampling algorithm, of which the first 2000 are discarded as burn-in draws and the remaining ones are kept for inference.[24]

4.3 In-Sample Results

Panel (a) of Figure 2 plots the posterior median, together with the 68% and 90% posterior credible intervals of the long-run growth rate of real GDP. This estimate is conditional on the entire sample and accounts for both filtering and parameter uncertainty. Several features of our estimate of long-run growth are worth noting. While the growth rate is stable between 3% and 4% during the first decades of the postwar period, a slowdown is clearly visible from around the late 1960's through the 1970's, consistent with the "productivity slowdown" (Nordhaus, 2004). The acceleration of the late 1990's and early 2000's associated with the productivity boom in the IT sector (Oliner and Sichel, 2000) is also visible. Thus, until the middle of the decade of the 2000's, our estimate conforms well to the generally accepted narrative about fluctuations in potential growth.[25] More recently, after peaking at about 3.5% in 2000, the median estimate of the long-run growth rate has fallen to about 2% in early 2015, a more substantial decline than the one observed after the productivity slowdown of the 1970's. Moreover, the slowdown appears to have happened gradually since the start

of around 0.4 percentage points in annualized terms, which is a fairly conservative prior.

[23] We provide further explanations and address robustness to the choice of priors in Appendix F.

[24] Thanks to the efficient state space representation discussed above, the improvements in the simulation smoother proposed by Bai and Wang (2015), and other computational improvements we implemented, the estimation is very fast. Convergence is achieved after only 1500 iterations, which take less than 20 minutes in MATLAB using an Intel 3.6 GHz computer with 16GB of DDR3 Ram.

[25] Appendix G provides a comparison of our estimate with the Congressional Budget Office (CBO) measure of potential growth, with some additional discussion.

of the 2000's, with most of the decline having occurred before the Great Recession.[26] Interestingly, a small rebound is visible at the end of the sample, but long-run growth stands far below its postwar average of 3.2%, with the 90% posterior credible interval ranging from 1.5% to 2.5%.

Panel (b) plots the time series of quarterly real GDP growth, together with the median posterior estimates of the common factor, aligned with the mean of real GDP growth. This plot highlights how the common factor captures the bulk of business-cycle frequency variation in output growth, while higher frequency, quarter-to-quarter variation is attributed to the idiosyncratic component. In the latter part of the sample, GDP growth is visibly below the factor, reflecting the decline in long-run growth.

The posterior estimate of the SV of the common factor is presented in Panel (c). It is clearly visible that volatility declines over the sample. The late 1940's and 1950's were extremely volatile, with a first large drop in volatility in the early 1960's. The Great Moderation is also clearly visible, with the average volatility pre-1985 being much larger than the average of the post-1985 sample. Notwithstanding the large increase in volatility during the Great Recession, our estimate of the common factor volatility since then remains consistent with the Great Moderation still being in place. This confirms the early evidence reported by Gadea-Rivas et al. (2014). It is clear from the figure that volatility spikes during recessions, a feature that brings our estimates close to the recent findings of Jurado et al. (2014) and Bloom (2014) relating to business-cycle uncertainty.[27] It appears that the random walk specification is flexible enough

[26]In principle, it is possible that our choice of modeling long-run GDP growth as a random walk is hard-wiring into our results the conclusion that the decline happened in a gradual way. In experiments with simulated data, presented in Appendix C, we show that if changes in long-run growth occur in the form of discrete breaks rather than evolving gradually, the (one-sided) filtered estimates will exhibit a discrete jump at the moment of the break. Instead, for US data the filtered estimates of the long-run growth component also decline in a gradual manner (see Figure A.1 in Appendix A).

[27]It is interesting to note that while in our model the innovations to the level of the common factor and its volatility are uncorrelated, the fact that increases in volatility are observed during recessions indicate the presence of negative correlation between the first and second moments, implying negative skewness in the distribution of the common factor. We believe a more explicit model of this feature

Figure 2: Trend, cycle and volatility: 1947-2015 (% Ann. Growth Rate)

(a) Posterior estimate of long-run growth

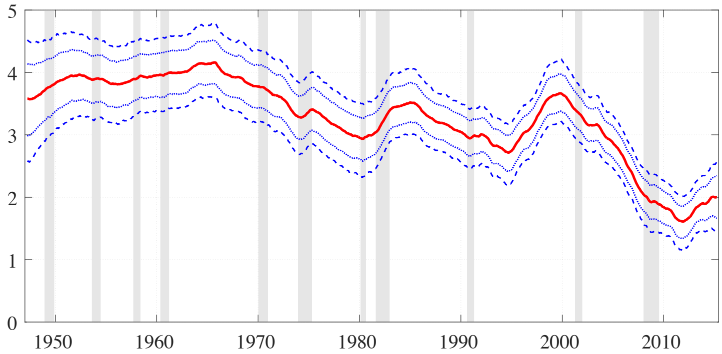

(b) Posterior estimate of common factor vs. actual GDP growth

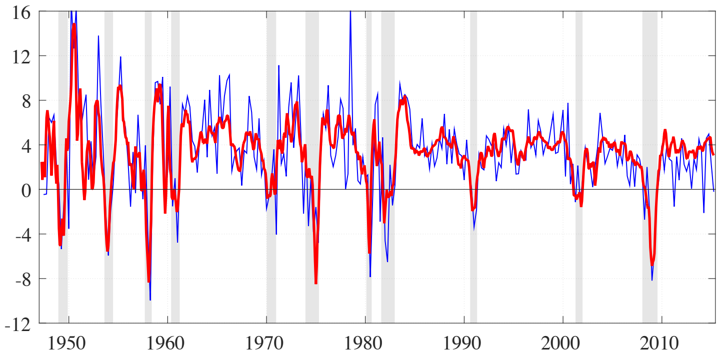

(c) Posterior estimate of common factor volatility

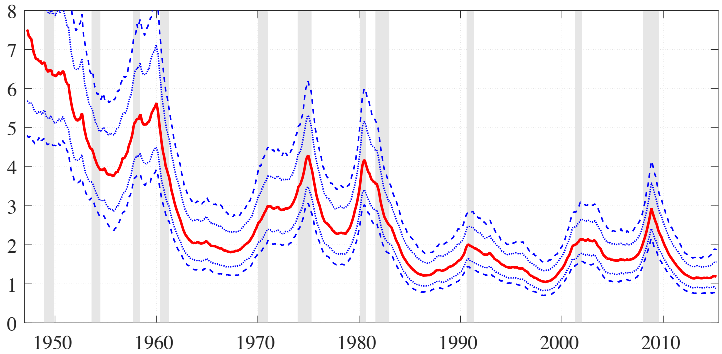

Note: Panel (a) displays the posterior median (solid red), together with the 68% and 90% (dotted and dashed blue) posterior credible intervals of long-run real GDP growth. Panel (b) plots actual real GDP growth (thin blue) against the posterior median estimate of the common factor, aligned with the mean of real GDP growth (thick red). Panel (c) presents the median (red), the 68% and the 90% (dotted and dashed blue) posterior credible intervals of the volatility of the common factor, i.e the square root of $\text{var}(f_t) = \sigma^2_{\varepsilon,t}(1-\phi_2)/[(1+\phi_2)((1-\phi_2)^2-\phi_1^2)]$. Shaded areas represent NBER recessions.

to capture cyclical changes in volatility as well as permanent phenomena such as the Great Moderation. Appendix A contains analogous charts for the volatilities of the idiosyncratic components of selected data series. Similar to the volatility of the common factor, many of the idiosyncratic volatilities present sharp increases during recessions.

The above results provide evidence that a significant decline in long-run US real GDP growth occurred over the last decade, and are consistent with a relatively gradual decline since the early 2000's. Our estimates show that the bulk of the slowdown from the elevated levels of growth at the turn of the century occurred before the Great Recession, which is consistent with the narrative of Fernald (2014) on the fading of the IT productivity boom. This recent decline is the largest movement in long-run growth observed in the postwar period.

4.4 Real-Time Results

As emphasized by Orphanides (2003), macroeconomic time series are heavily revised over time and in many cases these revisions contain valuable information that was not available at initial release. Therefore, it is important to assess, using the data available at each point in time, when the model detected the slowdown in long-run growth. For this purpose, we reconstruct our data set using vintages of data available from the Federal Reserve Bank of St. Louis ALFRED data base. Our aim is to replicate as closely as possible the situation of a decision-maker which would have applied our model in real time. We fix the start of our sample in 1947:Q1 and use an expanding out-of-sample window which starts on 11 January 2000 and ends on 30 June 2015. This is the longest possible window for which we are able to include the entire panel in Table 1 using fully real-time data. We then proceed by re-estimating the model each

is an important priority for future research.

day in which new data are released.[28]

Figure 3 looks at the model's real-time assessment of long-run growth at various points in time. Panel (a) plots the real-time estimate of current long-run growth, with 68% and 90% uncertainty bands. For comparison, the panel also shows the median response to the Philadelphia Fed Livingston Survey of Professional Forecasters (SPF) on the average growth rate for the next 10 years, and the estimate of long-run growth from a model with a constant intercept for GDP growth. The latter estimate is also updated as new information arrives, but weighs all points of the sample equally. Panel (b) displays vintages of the median long-run growth estimate, using information available up to July of each year. The locus traced by the end point of each vintage corresponds to the current real-time estimate of Panel (a).

The evolution of the baseline model's estimate of long-run growth when estimated in real time declines gradually from a peak of about 4% in early 2000 to around 2.5% just after the end of the Great Recession. From this time, the constant estimate shown in panel (a) is always outside of the 90% posterior bands. There is a sharp reassessment of long-run growth around July 2010, coinciding with the publication by the Bureau of Economic Analysis of the annual revisions to the National Accounts, which each year incorporate previously unavailable information for the previous three years. The revisions implied a substantial downgrade, in particular, to the growth of consumption in the first year of the recovery, from 2.5% to 1.6%, and instead allocated much of the growth in GDP during the recovery to inventory accumulation.[29] Reflecting the role of consumption as the most persistent and forward looking component of GDP,

[28] In a few cases new indicators were developed after January 2000. For example, the Markit Manufacturing PMI survey is currently one of the most timely and widely followed indicators, but it started being conducted in 2007. In those cases, we append to the panel, in real time, the vintages of the new indicators as soon sufficient history is available. In the example of the PMI, this is the case since mid-2012. By implication, the number of indicators in our data panel grows when new indicators appear. Full details about the construction of the vintage database are available in Appendix E.

[29] See Appendix I for additional figures on the National Accounts revisions during this period.

Figure 3: Long-Run GDP Growth Estimates in Real Time

(a) Evolution of the current assessment of long-run growth

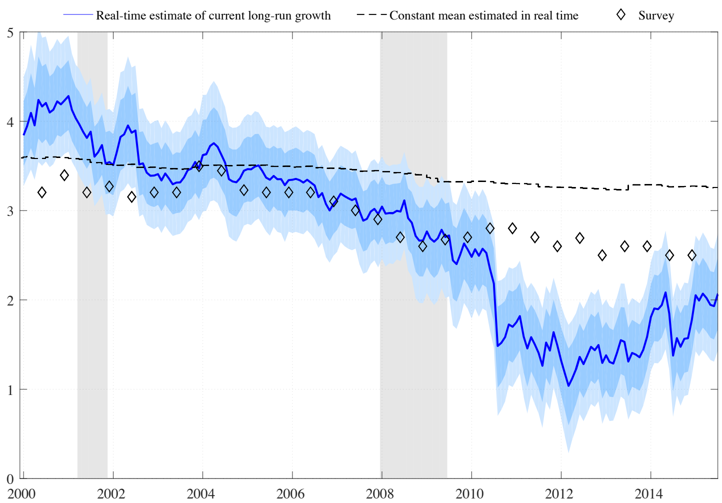

(b) Selected vintages of long-run growth estimates using real-time data

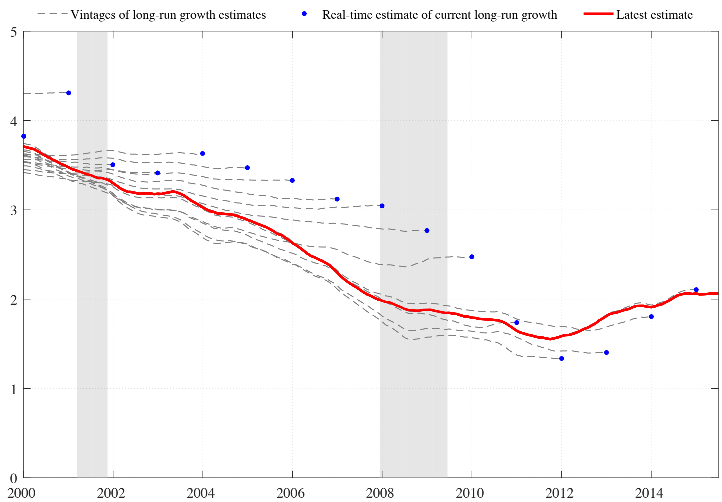

Note: The figure presents results from re-estimating the model using the vintage of data available at each point in time from January 2000 to March 2015. The start of the estimation sample is fixed at Q1:1947. Panel (a) plots the median real-time estimate of current long-run growth over time. This is the locus traced by the end points of all vintages. The blue shaded areas represent the 68th and 90th percentiles. The dashed line is the contemporaneous estimate of the historical average of real GDP growth. The diamonds are the median response to the Philadelphia Fed Livingston Survey of Professional Forecasters on the average growth rate for the next 10 years. Panel (b) displays the median estimate of long-run GDP growth for various vintages of data (dashed gray lines). The estimate of the latest vintage is shown in solid red. Gray shaded areas represent NBER recessions.

the estimate of long-run growth is downgraded sharply. Panel (b) shows how the 2010 revisions in fact trigger a re-interpretation of the years leading to the Great Recession. With the revised information, the bulk of the slowdown in long-run growth is now estimated to have occurred *before* the recession.[30] From 2010 onward, the model predicts a recovery that is extremely slow by historical standards. This is four years before the structural break test detected a statistically significant decline.[31] It is evident from the preceding discussion that revisions to past data by the BEA are an important source of changes to the long-run growth estimate in real time. Since the revision process is not modeled explicitly within the DFM, the in-sample results of Section 4.3 do not take into account the uncertainty stemming from future revisions. Interestingly, in the latest part of the sample, the estimate of long-run growth has recovered slightly to about 2% but this has been triggered by improvements in incoming data, rather than revisions to past vintages.

With regards to the SPF, it is noticeable that from 2003 to about 2010, the survey is remarkably similar to the model, but since then, the SPF forecast has continued to drift down only very slowly, standing at 2.5% as of mid-2015. It is noteworthy that, as pointed out by Stanley Fischer in the speech quoted in the introduction, during that period both private and institutional forecasters systematically overestimated growth.

4.5 Implications for Nowcasting GDP

The standard DFM with constant long-run growth and constant volatility has been successfully applied to produce current quarter nowcasts of GDP (see Banbura et al., 2010, for a survey). Using our real-time US database, we carefully evaluate whether

[30]Indeed, the (one-sided) filtered estimate based on the latest vintage, which ignores the effect of data revisions, displays a more gradual pattern of decline (see Figure A.1 in Appendix A).

[31]A simpler specification that does not use consumption to inform the trend would detect the decline in long-run growth one year later, with additional revisions to past GDP in July 2011.

the introduction of time-varying long-run growth and SV into the DFM framework also improves the performance of the model along this dimension. We find that over the evaluation window 2000-2015 the model is at least as accurate at point forecasting, and significantly better at density forecasting than the benchmark DFM. We find that most of the improvement in density forecasting comes from correctly assessing the center and the right tail of the distribution, implying that the time-invariant DFM is assigning excessive probability to a strong recovery. In an evaluation sub-sample spanning the post-recession period, the relative performance of both point and density forecasts improves substantially, coinciding with the significant downward revision of the model's assessment of long-run growth. In fact, ignoring the variation in long-run GDP growth would have resulted in being on average around 1 percentage point too optimistic from 2009 to 2015.[32]

To sum up, the addition of the time-varying components not only provides a tool for decision-makers to update their knowledge about the state of long-run growth in real time. It also brings about a substantial improvement in short-run forecasting performance when the trend is shifting, without worsening the forecasts when the latter is relatively stable. The proposed model therefore provides a robust and timely methodology to track GDP when long-run growth is uncertain.

4.6 Inspecting the Role of Data Set Size and Composition

In this paper we argue that the rich multivariate framework of a DFM will facilitate the extraction of the long-run growth component of GDP. The DFM will exploit the cross-sectional dimension, and not just the time series dimension in separating cycle from trend. It is interesting to quantify the advantage that the DFM provides over traditional trend-cycle decompositions, and to investigate the robustness of our main

[32] Appendix H provides the full details of the forecast evaluation exercise.

conclusions to alternative datasets of varying size and composition. In order to do so, we consider (1) a bivariate model with GDP and unemployment only (labeled "Okun"), (2) an intermediate model with GDP and the four additional variables often included in the construction of coincident indicators, see Mariano and Murasawa (2003) and Stock and Watson (1989) (labeled "MM03"), (3) our "Baseline" specification with 28 variables, and (4) an "Extended" model that uses disaggregated data for many of the headline series included in the baseline specification, totaling 155 variables.[33] Moreover, in order to investigate the gains associated with imposing additional structure to long-run GDP growth, for the last two specifications we also consider a version of the model that does not impose common long-run growth in GDP and consumption.

The top panel of Table 2 reports the mean point-estimates for each specification over selected subsamples.[34] In all cases, the results are consistent with a decline in the long-run growth rate in the last part of the sample. Quantitatively, most specifications are very close to the baseline, with the specifications that impose common long-run growth in GDP and consumption finding an earlier and sharper decline. The exception is the "Okun" specification which instead estimates a smaller increase in the mid 1990s as well as a larger decline in long-run growth in the past decade. It is noteworthy that the mean estimate of the extended specification is very close to that of the baseline.

The lower panel of Table 2 instead investigates the uncertainty around the mean estimates. The uncertainty around the long-run growth estimate declines as we move

[33] As we argue in Section 4.1, the introduction of a large number of disaggregated series, even if related to real activity, is likely to lead to model misspecification whenever the sectoral data are not contemporaneously related. For the extended specification, we consider a solution to this problem which allows to maintain the parsimonious one factor structure. By extending the model to include lags of the factor in the observation equation, each variable can display heterogeneous responses to the common factor, and correlation between idiosyncratic components is reduced. Given that the extended model is heavily parameterized, we follow D'Agostino et al. (2015) in choosing priors that shrink the model towards the contemporaneous-only specification, which is nested in the extended case. Full details and the composition of the data set and the changes to the estimation in case of the extended model are provided in Appendix J.

[34] See Figure J.1 in Appendix J for a comparison of the results of each alternative specification with the baseline results over the entire sample.

Table 2: Comparison of results for alternative data sets and specifications

	Okun	*MM03*	Baseline *GDP only*	Baseline ***GDP + C***	Extended *GDP only*	Extended *GDP + C*
Long-run growth						
1947-1972	3.9	3.5	3.6	**3.9**	3.7	3.9
1973-1995	3.2	3.4	3.1	**3.2**	3.2	3.2
1996-2007	2.6	3.2	3.0	**3.1**	3.0	3.1
2008-2015	1.5	2.5	2.5	**1.8**	2.1	1.7
End of Sample	1.2	2.4	2.3	**2.0**	2.0	2.1
Uncertainty: Long run						
Filtered	0.84	0.62	0.63	**0.53**	0.70	0.62
Smoothed	0.44	0.36	0.35	**0.32**	0.41	0.39
Uncertainty: Cycle						
Filtered	2.15	1.48	0.79	**0.78**	0.23	0.18
Smoothed	1.98	1.32	0.62	**0.61**	0.25	0.17

Notes: Each column presents the estimation results corresponding to the alternative models (data sets) considered in this section. The upper panel displays the posterior means of the long-run growth rate of real GDP, over selected subsamples. In the lower panel, the posterior uncertainty corresponding to both the long-run growth rate of real GDP, as well as the common factor are displayed. The uncertainty is calculated as an average over the entire sample.

from the bivariate to the multivariate specifications, with most of the reduction happening once a handful of variables are included. On the other hand, when the panel is extended to include a large number of disaggregated series, the uncertainty increases.[35] While including a few key series, such as the ones in the specification of Mariano and Murasawa (2003) seems to already achieve the bulk of the reduction in uncertainty, it should be taken into account that those variables are available only with a relatively long publication lag, and subject to considerable revisions over time. Our proposed strategy of using an intermediate number of indicators, including the more timely and accurate surveys, is likely to lead to more satisfactory results in a real-time setting. Furthermore, the inclusion of the surveys is helpful in identifying the long-run growth

[35]We conjecture that as many more variables are added, the fit of the common factor to the cyclical component of GDP worsens. As a consequence, some cyclical variation of GDP spills over to the estimate of the long-run component. The uncertainty around the common factor, on the other hand, continues to decline.

rate, as those variables do not display a time-varying long-run mean by construction.

Overall this exercise highlights that the finding of a substantial decline in the long-run growth rate is confirmed across different specifications that use data sets of varying size and composition. The baseline specification, which uses an intermediate number of series including both hard data and surveys, leads to the lowest uncertainty around the long-run growth estimate, supporting the baseline choice of data set size and composition proposed in Section 4.1. Our results have important implications for trend-cycle decompositions of output, which usually include only a few cyclical indicators, generally inflation or variables that are direct inputs to the production function (see e.g. Gordon, 2014a or Reifschneider et al., 2013). As we show, greater precision of the trend component can be achieved by exploiting the common cyclical features of additional macroeconomic variables.[36]

5 Decomposing Movements in Long-Run Growth

In this section, we show how our model can be used to decompose the long-run growth rate of output into long-run movements in labor productivity and labor input. By doing this, we exploit the ability of the model to filter away cyclical variation and idiosyncratic noise and obtain clean estimates of underlying long-run trends. We see this exercise as a step towards giving an economically more meaningful interpretation to the movements in long-run real GDP growth detected by our model.

GDP growth is by identity the sum of growth in output per hour and growth in total hours worked. It is therefore possible to split the long-run growth trend in our model into two orthogonal components such that this identity is satisfied in the long

[36]Basistha and Startz (2008) make a similar point, arguing that the inclusion of indicators that are informative about common cycles can help reduce the uncertainty around Kalman filter estimates of the long-run rate of unemployment (NAIRU).

run. Here we make use of our flexible definition of $\mathbf{c}_t$ in equation (2). In particular, ordering the growth rates of real GDP, real consumption and total hours as the first three variables in $\mathbf{y}_t$, we define

$$\mathbf{a}_t = \begin{bmatrix} z_t \\ h_t \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 1 & 1 \\ 1 & 1 \\ 0 & 1 \end{bmatrix}, \tag{10}$$

so that the model is specified with two time-varying components, the first of which loads output and consumption but not hours, and the second loads all three series. The first component is then by construction the long-run growth rate of labor productivity, while the second one captures low-frequency movements in labor input independent of productivity.[37] Given the relation in (10), the two components add up to the time-varying intercept in the baseline specification, i.e. $g_t = z_t + h_t$.[38] It follows from standard growth theory that our estimate of the long-run growth rate of labor productivity will capture both technological factors and other factors, such as capital deepening and labor quality.[39]

Figure 4 presents the results of the decomposition exercise for the US. Panel (a) plots the median posterior estimate of long-run real GDP growth and its labor productivity and total hours components. The posterior bands for long-run real GDP

[37] z_t and h_t jointly follow random walks with diagonal covariance matrix as defined by equation (7). Restricting the covariance matrix is not necessary for estimation, but imposing it allows us to interpret the innovations to the trends as exogenous shocks to the long-run growth rates of the variables. The hours trend is therefore interpreted as those low-frequency movements in hours which are uncorrelated with labor productivity. Allowing for a full covariance matrix would yield trends that are linear combinations of the current ones, but would lack a clear economic interpretation.

[38] Since z_t and h_t are independent and add up to g_t, we set the prior on the scale of their variances to half of the one set in Section 4.2 on g_t. In addition, note that the cyclical movement in labor productivity is given by $(1 - \lambda_3) f_t$.

[39] Further decomposing z_t into technology and non-technology movements requires additional information to separately identify these components. One possibility, which we explore in Appendix K, is to use an independent measure of TFP to isolate technological factors. Note, however, that reliable data on capital input, labor quality, or estimates of TFP are not available in real time, making the focus on long-run labor productivity more appealing in a real-time setting.

Figure 4: Decomposition of Long-run US Output Growth

(a) Posterior median estimates of decomposition

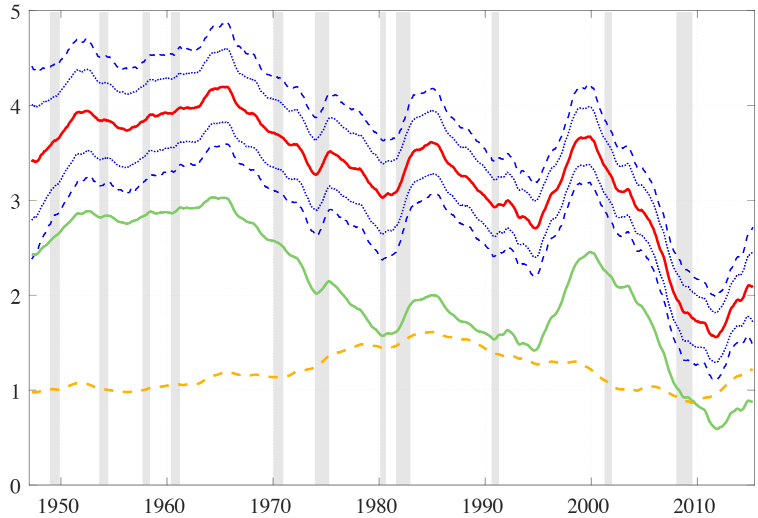

(b) Filtered estimates of long-run growth components

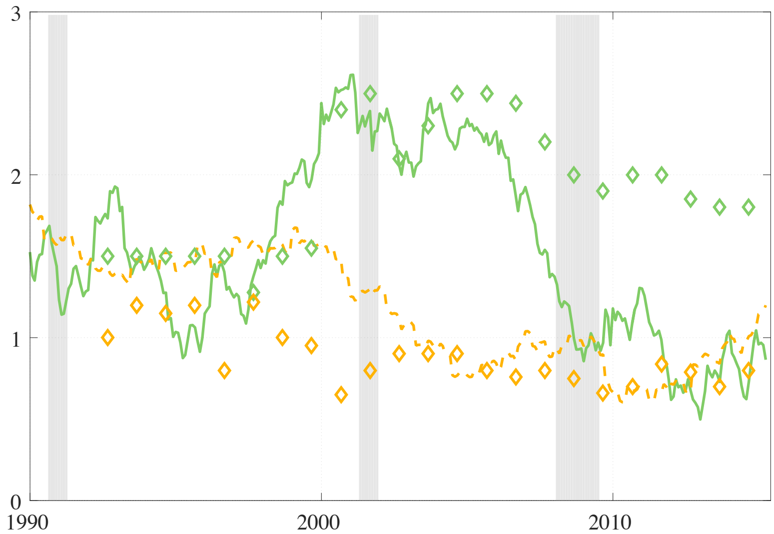

Note: Panel (a) plots the posterior median (solid red), together with the 68% and 90% (doted and dashed blue) posterior credible intervals of long-run GDP growth and the posterior median of both long-run labor productivity growth and long-run total hours growth (solid green and dashed orange). Panel (b) plots the filtered estimates of these two components, i.e. $\hat{z}_{t|t}$ and $\hat{h}_{t|t}$, since 1990. For comparison, the corresponding forecasts from the SPF are plotted. The SPF forecast for total hours is obtained as the difference between the forecasts for real GDP and labor productivity.

growth are included. The time series evolution conforms very closely to the narrative of Fernald (2014), with a pronounced boom in labor productivity in the mid-1990's and a subsequent fall in the 2000's clearly visible. The decline in the 2000's is relatively sudden while the 1970's slowdown appears as a more gradual phenomenon starting in the late 1960's. Furthermore, the results reveal that during the 1970's and 1980's the impact of the productivity slowdown on output growth was partly masked by a secular increase in hours, probably reflecting increases in the working-age population as well as labor force participation (see e.g. Goldin, 2006). Focusing on the period since 2000, labor productivity accounts for almost the entire decline.[40] This contrasts explanations by which slow labor force growth has been a drag on GDP growth. When taking away the cyclical component of hours and focusing solely on its long-run component, the contribution of hours has, if anything, accelerated since the Great Recession. Panel (b) presents the filtered estimates of the two components, i.e. the output of the Kalman Filter which uses data only up to each point in time. For comparison, the corresponding SPF forecasts are included. Most notably, this plot reveals that starting around 2005 a relatively sharp revision to labor productivity drives the decline in long-run output growth.[41] Interestingly, the professional forecasters have been very slow in incorporating the productivity slowdown into their long-run forecasts. This delay explains their persistent overestimation of GDP growth since the recession.

It is interesting to compare the results of our decomposition exercise to similar approaches in the literature, in particular Gordon (2010, 2014a) and Reifschneider et al. (2013). Like us, they specify a state space model with a common cyclical component and use the 'output identity' to decompose the long-run growth rate of GDP into

[40]In Appendix K we extend the analysis to decompose the labor productivity trend into long-run TFP and non-technological forces. We find that TFP accounts for virtually all of the slowdown.

[41]In an additional figure, provided in Appendix A, we plot 5,000 draws from the joint posterior distribution of the variances of the innovations to the labor productivity and hours components. This analysis confirms the conclusion from the discussion here that changes in labor productivity, rather than in labor input, are the key driver of low frequency movements in real GDP growth.

underlying drivers. A key difference resides in the Bayesian estimation of the model, which enables us to impose a conservative prior on the variance of the long-run growth component that helps avoiding over-fitting the data. Furthermore, the inclusion of SV in the cyclical component helps to prevent unusually large cyclical movements from contaminating the long-run estimate. Another important difference is that we use a larger amount of information, including key cyclical indicators like industrial production, sales, and business surveys, which are generally not included in a production function approach. This allows us to retrieve a timely and precise estimate of the cyclical component and, as a consequence, to reduce the uncertainty that is inherent to any trend-cycle decomposition of the data, as discussed in Section 4.6. As a result, we obtain a substantially less pessimistic estimate of the long-run growth of GDP than these studies in the latest part of the sample. For instance, Gordon (2014a) reports a long-run GDP growth estimate below 1% for the end of the sample, whereas our median estimate stands at around 2%.[42]

5.1 International Evidence

To gain an international perspective on our results, we estimate the DFM for the other G7 economies and perform the decomposition exercise for each of them.[43] The median posterior estimates of the labor productivity and labor input trends are displayed in Figure 5. Labor productivity, displayed in Panel (a), plays again the key role

[42] The results for a bivariate model of GDP and unemployment, which we have discussed in Section 4.6 show that the current long-run growth estimate is 1.2%, close to Gordon (2014a).

[43] Details on the specific data series used for each country are available in Appendix E. For hours, we again follow the methodology of Ohanian and Raffo (2012). In the particular case of the UK, the quarterly series for hours displays a drastic change in its stochastic properties in the early 1990's owing to a methodological change in the construction by the ONS, as confirmed by the ONS LFS manual. We address this issue by using directly the annual series from the TED, which requires an appropriate extension of equation (9) to annual variables (see Banbura et al. 2012). To avoid weak identification of h_t for the UK, we truncate our prior on its variance to discard values which are larger than twice the maximum posterior draw of the case of the other countries.

in determining movements in long-run growth. In the Western European economies and Japan, the elevated growth rates of labor productivity prior to the 1970's reflect the rebuilding of the capital stock from the destruction from World War II, and ended as these economies converged towards US levels of output per capita. The labor productivity profile of Canada broadly follows that of the US, with a slowdown in the 1970's and a temporary mild boom during the late 1990's. Interestingly, this acceleration in the 1990's did not occur in Western Europe and Japan.[44] The UK displays a decline in labor productivity similar to the US. This "productivity puzzle" has been debated extensively in the UK (see e.g. Pessoa and Van Reenen, 2014). It is interesting to note that the two countries which experienced a more severe financial crisis, the US and the UK, appear to be the ones with greatest declines in productivity since the early 2000's, similar to the evidence documented in Reinhart and Rogoff (2009).

Panel (b) displays the movements in long-run hours worked identified by equation (10). The contribution of this component to overall long-run output growth varies considerably across countries. However, within each country it is more stable over time than the productivity component, which is in line with our findings for the US. Indeed, the extracted long-run trend in total hours includes various potentially offsetting forces that can lead to changes in long-run output growth. In any case, the results of our decomposition exercise indicate that after using the DFM to remove business-cycle variation in hours and output, the decline in long-run GDP growth that has been observed in the advanced economies since the early 2000's is entirely accounted for by a decline in the labor productivity trend. Finally, it is interesting to note that for the countries in the sample long-run productivity growth appears to converge in the cross section, while there is no evidence of convergence in the long-run growth of hours.[45]

[44]On the lost decade in Japan, see Hayashi and Prescott (2002). Gordon (2004) examines the absence of the IT boom in Europe.

[45]Similar evidence for emerging economies has been recently presented by Pritchett and Summers (2014). Their evidence refers to convergence of overall GDP growth rates, whereas ours indicates that

Figure 5: Decomposition for Other Advanced Economies

(a) Long-run Labor Productivity

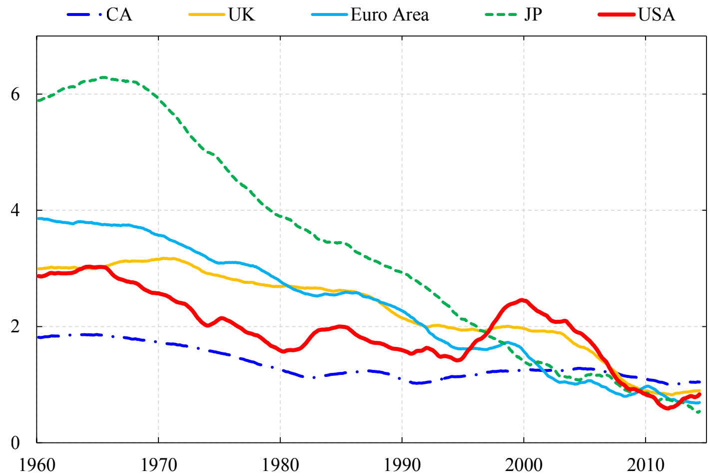

(b) Long-run Labor Input

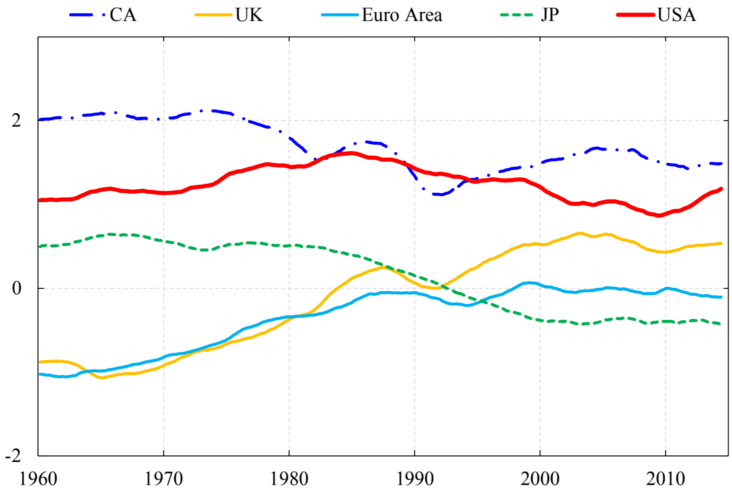

Note: Panel (a) displays the posterior median of long-run labor productivity across advanced economies. Panel (b) plots the corresponding estimates of long-run total hours worked. In both panels, 'Euro Area' represents a weighted average of Germany, Italy and France.

6 Concluding Remarks

The sluggish recovery from the Great Recession has raised the question whether the long-run growth rate of US real GDP is now lower than it has been on average over the postwar period. We have presented a dynamic factor model that allows for both changes in long-run GDP growth and stochastic volatility. Estimating the model with Bayesian methods, we provide evidence that long-run growth of US GDP displays a gradual decline after the turn of the century, moving from its peak of 3.5% to about 2% in 2015. Using real-time vintages of data we demonstrate the model's ability to track GDP in a timely and reliable manner. By the summer of 2010 the model would have concluded that a significant decline in long-run growth was behind the slow recovery, therefore substantially improving the real-time tracking of GDP by explicitly taking into account the uncertainty surrounding long-run growth. Finally, we discuss the drivers of movements in long-run output growth through the lens of our model by decomposing it into the long-run growth rates of labor productivity and labor input. Using data for both the US and other advanced economies our model points to a global slowdown in labor productivity as the main driver of weak growth in recent years, extending the narrative of Fernald (2014) to other economies. Studying the deep causes of the secular decline in growth is an important priority for future research.

convergence in productivity growth appears to be the dominant source of convergence.

References

Alvarez, R., Camacho, M., and Perez-Quiros, G. (2012). Finite sample performance of small versus large scale dynamic factor models. CEPR Discussion Papers 8867, C.E.P.R. Discussion Papers.

Auerbach, A. J. (2011). Long-term fiscal sustainability in major economies. BIS Working Papers 361, Bank for International Settlements.

Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1):47–68.

Bai, J. and Wang, P. (2015). Identification and Bayesian Estimation of Dynamic Factor Models. *Journal of Business & Economic Statistics*, 33(2):221–240.

Banbura, M., Giannone, D., Modugno, M., and Reichlin, L. (2012). Now-casting and the real-time data flow. Working Papers ECARES 2012-026, ULB.

Banbura, M., Giannone, D., and Reichlin, L. (2010). Nowcasting. CEPR Discussion Papers 7883, C.E.P.R. Discussion Papers.

Banbura, M. and Modugno, M. (2014). Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data. *Journal of Applied Econometrics*, 29(1):133–160.

Basistha, A. and Startz, R. (2008). Measuring the NAIRU with Reduced Uncertainty: A Multiple-Indicator Common-Cycle Approach. *The Review of Economics and Statistics*, 90(4):805–811.

Benati, L. (2007). Drift and breaks in labor productivity. *Journal of Economic Dynamics and Control*, 31(8):2847–2877.

Beveridge, S. and Nelson, C. R. (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the business cycle. *Journal of Monetary Economics*, 7(2):151–174.

Bloom, N. (2014). Fluctuations in Uncertainty. *Journal of Economic Perspectives*, 28(2):153–76.

Boivin, J. and Ng, S. (2006). Are more data always better for factor analysis? *Journal of Econometrics*, 132(1):169–194.

Clark, P. K. (1987). The Cyclical Component of U.S. Economic Activity. *The Quarterly Journal of Economics*, 102(4):797–814.

Cochrane, J. H. (1994). Permanent and Transitory Components of GNP and Stock Prices. *The Quarterly Journal of Economics*, 109(1):241–65.

Cogley, T. (2005). How fast can the new economy grow? A Bayesian analysis of the evolution of trend growth. *Journal of Macroeconomics*, 27(2):179–207.

Cogley, T. and Sargent, T. J. (2005). Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics*, 8(2):262–302.

Creal, D., Koopman, S. J., and Zivot, E. (2010). Extracting a robust US business cycle using a time-varying multivariate model-based bandpass filter. *Journal of Applied Econometrics*, 25(4):695–719.

D'Agostino, A., Giannone, D., Lenza, M., and Modugno, M. (2015). Nowcasting business cycles: A bayesian approach to dynamic heterogeneous factor models. Technical report.

Del Negro, M. and Otrok, C. (2008). Dynamic factor models with time-varying parameters:measuring changes in international business cycles. NYFed Staff Rep. 326.

Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods: Second Edition*. Oxford University Press.

Edge, R. M., Laubach, T., and Williams, J. C. (2007). Learning and shifts in long-run productivity growth. *Journal of Monetary Economics*, 54(8):2421–2438.

Evans, M. D. D. (2005). Where Are We Now? Real-Time Estimates of the Macroeconomy. *International Journal of Central Banking*, 1(2).

Fernald, J. (2014). Productivity and potential output before, during, and after the great recession. *NBER Macroeconomics Annual 2014*, 29.

Fernald, J. G. and Jones, C. I. (2014). The Future of US Economic Growth. *American Economic Review*, 104(5):44–49.

Fleischman, C. A. and Roberts, J. M. (2011). From many series, one cycle: improved estimates of the business cycle from a multivariate unobserved components model. Finance and Economics Discussion Series 2011-46, Fed Board of Governors.

Forni, M., Giannone, D., Lippi, M., and Reichlin, L. (2009). Opening The Black Box: Structural Factor Models With Large Cross Sections. *Econometric Theory*, 25(05):1319–1347.

Forni, M., Hallin, M., Lippi, M., and Reichlin, L. (2003). Do financial variables help forecasting inflation and real activity in the euro area? *Journal of Monetary Economics*, 50(6):1243–1255.

Gadea-Rivas, M. D., Gomez-Loscos, A., and Perez-Quiros, G. (2014). The Two Greatest. Great Recession vs. Great Moderation. WP Series 1423, Bank of Spain.

Geweke, J. (1977). The dynamic factor analysis of economic time series. In *Latent Variables in Socio-Economic Models*. North-Holland.

Giannone, D., Reichlin, L., and Sala, L. (2006). Vars, common factors and the empirical validation of equilibrium business cycle models. *Journal of Econometrics*, 132(1):257–279.

Giannone, D., Reichlin, L., and Small, D. (2005). Nowcasting gdp and inflation: the real-time informational content of macroeconomic data releases. Technical report.

Giannone, D., Reichlin, L., and Small, D. (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55(4):665–676.

Goldin, C. (2006). The quiet revolution that transformed women's employment, education, and family. *AEA Papers and Proceedings*, May 2006.

Gordon, R. J. (2004). Why was europe left at the station when america's productivity locomotive departed? Working Paper 10661, National Bureau of Economic Research.

Gordon, R. J. (2010). Okun's Law and Productivity Innovations. *American Economic Review*, 100(2):11–15.

Gordon, R. J. (2014a). A new method of estimating potential real gdp growth: Implications for the labor market and the debt/gdp ratio. NBER Working Paper 20423.

Gordon, R. J. (2014b). The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections. NBER Working Papers 19895.

Greenwood, J., Hercowitz, Z., and Krusell, P. (1997). Long-Run Implications of Investment-Specific Technological Change. *American Economic Review*, 87(3):342–62.

Hansen, B. E. (1992). Testing for parameter instability in linear models. *Journal of Policy Modeling*, 14(4):517–533.

Harvey, A. C. (1985). Trends and Cycles in Macroeconomic Time Series. *Journal of Business & Economic Statistics*, 3(3):216–27.

Harvey, A. C. and Stock, J. H. (1988). Continuous time autoregressive models with common stochastic trends. *Journal of Economic Dynamics and Control*, 12(2-3):365–384.

Hayashi, F. and Prescott, E. C. (2002). The 1990s in japan: A lost decade. *Review of Economic Dynamics*, 5(1):206 – 235.

Hodrick, R. J. and Prescott, E. C. (1997). Postwar U.S. Business Cycles: An Empirical Investigation. *Journal of Money, Credit and Banking*, 29(1):1–16.

Jacquier, E., Polson, N. G., and Rossi, P. E. (2002). Bayesian Analysis of Stochastic Volatility Models. *Journal of Business & Economic Statistics*, 20(1):69–87.

Jorgenson, D. W., Ho, M. S., and Stiroh, K. J. (2006). Potential Growth of the U.S. Economy: Will the Productivity Resurgence Continue? *Business Economics*, 41(1):7–16.

Jurado, K., Ludvigson, S. C., and Ng, S. (2014). Measuring uncertainty. *American Economic Review, Forthcoming.*

Kim, C.-J. and Nelson, C. R. (1999a). Has The U.S. Economy Become More Stable? A Bayesian Approach Based On A Markov-Switching Model Of The Business Cycle. *The Review of Economics and Statistics*, 81(4):608–616.

Kim, C.-J. and Nelson, C. R. (1999b). *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*. The MIT Press.

Kim, S., Shephard, N., and Chib, S. (1998). Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Studies*, 65(3):361–93.

Luciani, M. and Ricci, L. (2014). Nowcasting Norway. *International Journal of Central Banking, Forthcoming.*

Luo, S. and Startz, R. (2014). Is it one break or on going permanent shocks that explains u.s. real gdp? *Journal of Monetary Economics*, 66:155–163.

Marcellino, M., Porqueddu, M., and Venditti, F. (2014). Short-term gdp forecasting with a mixed frequency dynamic factor model with stochastic volatility. *Journal of Business & Economic Statistics, Forthcoming.*

Mariano, R. S. and Murasawa, Y. (2003). A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics*, 18(4):427–443.

McConnell, M. M. and Perez-Quiros, G. (2000). Output fluctuations in the united states: What has changed since the early 1980's? *American Economic Review*, 90(5):1464–1476.

Morley, J. C., Nelson, C., and Zivot, E. (2003). Why are the beveridge-nelson and unobserved-components decompositions of gdp so different? *The Review of Economics and Statistics*, 85(2):235–243.

Nelson, C. R. and Plosser, C. I. (1982). Trends and random walks in macroeconomic time series: Some evidence and implications. *Journal of Monetary Economics*, 10(2):139–162.

Nordhaus, W. (2004). Retrospective on the 1970s Productivity Slowdown. Nber working papers, National Bureau of Economic Research, Inc.

Nyblom, J. (1989). Testing for the constancy of parameters over time. *Journal of the American Statistical Association*, 84(405):223–230.

Ohanian, L. E. and Raffo, A. (2012). Aggregate hours worked in OECD countries: New measurement and implications for business cycles. *Journal of Monetary Economics*, 59(1):40–56.

Oliner, S. D. and Sichel, D. E. (2000). The Resurgence of Growth in the Late 1990s: Is Information Technology the Story? *Journal of Economic Perspectives*, 14(4):3–22.

Orphanides, A. (2003). The quest for prosperity without inflation. *Journal of Monetary Economics*, 50(3):633–663.

Pain, N., Lewis, C., Dang, T.-T., Jin, Y., and Richardson, P. (2014). OECD Forecasts During and After the Financial Crisis: A Post Mortem. OECD Economics Department Working Papers 1107, OECD Publishing.

Perron, P. and Wada, T. (2009). Let's take a break: trends and cycles in us real gdp. *Journal of Monetary Economics*, 56:749–765.

Pessoa, J. P. and Van Reenen, J. (2014). The uk productivity and jobs puzzle: Does the answer lie in wage flexibility? *The Economic Journal*, 124(576):433–452.

Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72(3):821–852.

Pritchett, L. and Summers, L. H. (2014). Asiaphoria Meets Regression to the Mean. NBER Working Papers 20573, National Bureau of Economic Research, Inc.

Reifschneider, D., Wascher, W., and Wilcox, D. (2013). Aggregate supply in the United States: recent developments and implications for the conduct of monetary policy. Finance and Economics Discussion Series 2013-77, Board of Governors of the Fed.

Reinhart, C. M. and Rogoff, K. S. (2009). The Aftermath of Financial Crises. *American Economic Review*, 99(2):466–72.

Sargent, T. J. and Sims, C. A. (1977). Business cycle modeling without pretending to have too much a priori economic theory. Working Papers 55, Federal Reserve Bank of Minneapolis.

Sims, C. A. (2012). Comment to stock and watson (2012). *Brookings Papers on Economic Activity*, Spring:81–156.

Stock, J. and Watson, M. (2012). Disentangling the channels of the 2007-09 recession. *Brookings Papers on Economic Activity*, Spring:81–156.

Stock, J. H. and Watson, M. W. (1989). New Indexes of Coincident and Leading Economic Indicators. In *NBER Macroeconomics Annual 1989, Volume 4*, NBER Chapters, pages 351–409. National Bureau of Economic Research, Inc.

Stock, J. H. and Watson, M. W. (1999). Business cycle fluctuations in us macroeconomic time series. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1 of *Handbook of Macroeconomics*, chapter 1, pages 3–64. Elsevier.

Stock, J. H. and Watson, M. W. (2002). Forecasting Using Principal Components From a Large Number of Predictors. *Journal of the American Statistical Association*, 97:1167–1179.

Stock, J. H. and Watson, M. W. (2003). Has the business cycle changed and why? *NBER Macroeconomics Annual 2002*, 17:159–230.

Summers, L. (2014). Secular stagnation. IMF Economic Forum: Policy Responses to Crises, speech delivered at the IMF Annual Research Conference, November 8th.

Whelan, K. (2003). A Two-Sector Approach to Modeling U.S. NIPA Data. *Journal of Money, Credit and Banking*, 35(4):627–56.