Is economic growth permanently lower?*

Juan Antolin-Diaz[(1)], Thomas Drechsel[(2)] and Ivan Petrella[(3)]

The slow pace of recovery from the 2007-2009 recession has recently prompted questions about whether the long-run growth rate of the US economy (and other industrialised economies) is lower now than it has been on average over the past few decades. Indeed, for the last five years, forecasts of US and global real GDP growth have been persistently biased upwards. We extend a state-of-the-art econometric model which tracks economic activity to make it robust to changes in long-run growth. This allows us to uncover a substantial decline in long-run growth in the US and other industrialised economies.

Real activity forecasts have been persistently biased upwards in recent years (Figure 1), both in the US and in other industrialised economies. Every year since the recovery from the Financial Crisis started economic projections have been progressively revised downwards as the year advanced. What accounts for these persistent disappointments? In this note we use a state-of-the-art econometric model to provide evidence of a persistent decline in economic growth. Failure to account for this decline in long-run growth may be behind the persistent disappointments in forecasting economic activity.

Figure 1. US GDP: FOMC projections and actual outturn

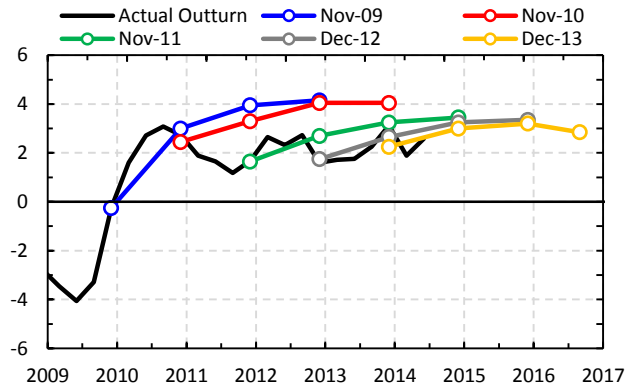

Source: FOMC

Evidence of breaks in the mean of GDP

As means of preliminary evidence of a slowdown in GDP growth, we test for multiple breaks in the mean of GDP growth. We find that there is evidence at the 5% level in favour of at least one break. The most likely break is in the second quarter of 2000, while the second most likely break, which is not significant, is estimated to have occurred in the second quarter of 1973.

These results are similar to those of Luo and Startz (2014), who use Bayesian model averaging to calculate the posterior probability of a single break and find the most likely break date in 2006:Q1. They note that if the sample were restricted to exclude the decade of the 2000's, a break date around 1973:Q1 would be the most likely. This is also in line with the analysis of the labour productivity series by Fernald (2014), who finds

Figure 2. Break in US GDP growth as detected by Bai-Perron Test

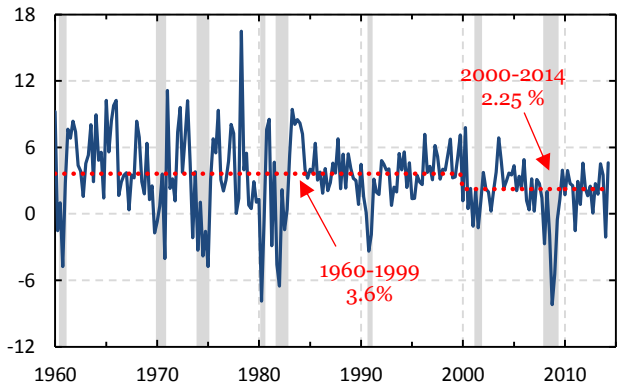

Figure 3. Real-Time Result of the Bai-Perron Test

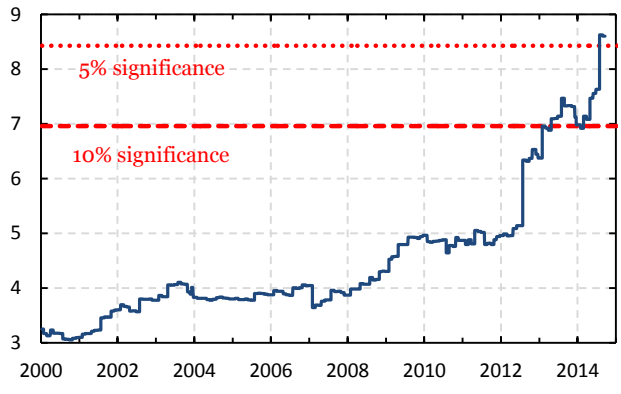

* This note is an abridged version of the technical working paper "Following the Trend: Tracking GDP when Long-Run Growth is Uncertain", available for download here. (1) Fulcrum Asset management; (2) LSE; (3) Birkbeck and CEPR.

evidence for three breaks in the mean growth rate of productivity. First, a productivity slowdown in 1973:Q2, second, a speedup around 1995:Q3, and finally a second slowdown in the early 2000's. Given the somewhat different results obtained with different samples and methods, the precise number and timing of breaks remains unclear to us. However, we note that there is substantial evidence for at least one break in the mean of GDP in the post-War sample, most likely in the first half of the decade of the 2000's.

It is noteworthy that the early 2000's break detected by the test only became significant at standard levels with the recent vintages of National Accounts data, as displayed in Figure 3. If the test is correct and the break happened at the beginning of the decade, this means that the break was not detected until almost fifteen years later. This highlights the importance of an econometric framework capable of detecting changes in long-run growth in a timely and accurate manner.

Tracking GDP when long-run growth is uncertain

Orphanides (2003) emphasized that real-time misperceptions about the long-run growth of the economy can have a large role in monetary policy mistakes, so the possibility of significant instabilities in the mean of real output calls for a robust framework capable of producing timely assessments of short- and long-run GDP growth.

Since the seminal contribution of Giannone et al. (2008) the Dynamic Factor Model (DFM) has become one of the most important tools to produce real-time tracking estimates (also known as "nowcasts") of GDP[1]. DFMs capture the idea that a small number of unobserved factors drive 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. Giannone, Reichlin and Small (2008) and Banbura et al. (2012) have applied the DFM framework to the problem of nowcasting GDP; that is, obtaining early estimates of quarterly GDP growth by exploiting more timely monthly indicators and the ability of the model to incorporate quarterly and monthly time series within a unified framework.

To take seriously the possibility that long-run growth might be shifting over time, we introduce two new features into an otherwise standard DFM of real activity data: a time-varying long-run growth component for GDP, and stochastic volatility (SV) in the innovations to both factors and idiosyncratic components. We apply the model to a broad panel with 26 variables of US real activity data, and calculate the probability distribution of long-run GDP growth (Figure 4).

Figure 4. Evolution of US GDP long-run growth

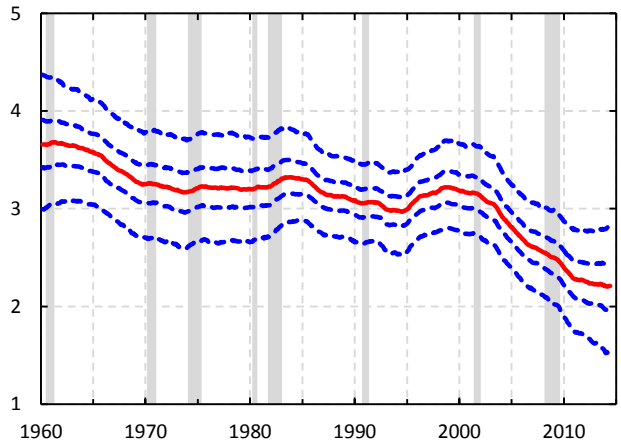

Note: This figure plots the posterior median (solid red), together with the 68% and 90% (dashed blue) posterior credible intervals of our long-run GDP growth estimate. Shaded areas represent NBER recessions.

An initial slowdown is visible around the late 1960's, close to the well-known 1973 "productivity slowdown". The acceleration of the late 1990's and early 2000's associated with the productivity boom in the IT sector is also clearly visible. Thus, until the middle of the decade of the 2000's, our estimate conforms well to the generally accepted narrative about fluctuations in

[1] Recently these authors have made their nowcasts available to financial market participants. See http://www.now-casting.com

potential growth. It must be noted, however, that according to our estimates until the most recent part of the sample, the historical average of 3.15% is always contained within the 90% credible interval. This is consistent with the fact that the break tests do not find significant breaks during 1973 or the mid-1990's. Finally, from its peak of about 3.25% in late 1998 to its level as of June 2014, 2.25%, the median estimate of the trend has declined by one percentage point, a more substantial decline than the one observed after the original "productivity slowdown" of the 1970's. Moreover, the decline appears to have happened gradually since the start of the 2000's, with the reductions in trend growth clustered around two episodes: one around the middle of the decade, but before the Financial Crisis, and a second after the recession, at the beginning of the subsequent recovery.

Real-Time Evidence of a Slowdown in US GDP

As is well known, macroeconomic time series are revised (sometimes heavily) over time, and in many cases these revisions contain valuable information that was not available at initial release. Therefore, it is possible that our results are only apparent using the current vintage of data, and our model would not have been able to detect the slowdown as it happened. To address this concern, we reconstruct our dataset at each point in time, using vintages of data available from the Federal Reserve Bank of St. Louis ALFRED database. We document how, by the summer of 2011 the model would have concluded that a significant decline in long-run growth was behind the slow recovery. That is, three years before the break test gave a conclusive answer.

Figure 5. Evolution of US GDP long-run growth in Real Time

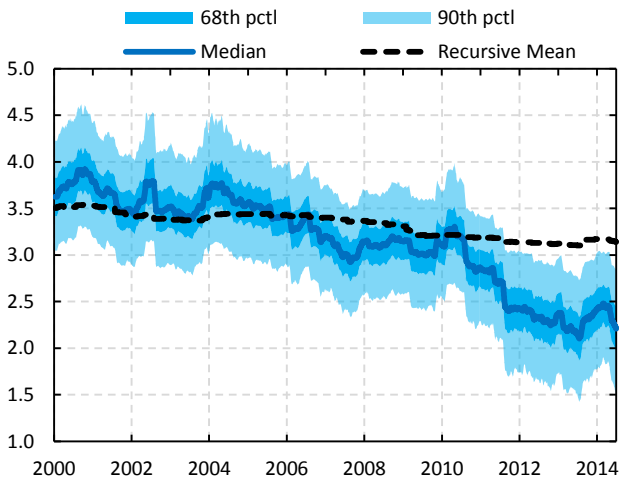

Note: The shaded areas represent the 68th and 90th percentile, together with the median of the posterior credible interval of the current value of long-run GDP growth, re-estimated daily from January 2000 to September 2014 using the vintage of data available at each point in time. The blacked dashed line the contemporaneous estimate of the historical average of GDP growth rate with the start of the sample fixed at Q1:1960.

Long-run growth fluctuations in the G7

So far we have focused our discussion on the US economy, where long-run growth had been remarkably stable until the turn of century. The analysis of data from other industrialized economies indicates that the post-War stability of the US long-run growth rate is an exception rather than the rule.Figure 6 plots the estimate of long-run growth resulting from re-estimating our model with data for each of the other G7 economies. Large fluctuations in trend growth are apparent for most of the countries: during the 1960's, Germany, France and Italy were growing by 4% on a sustained basis, while Japan was growing by about 7%. These high growth rates were a consequence of the need to rebuild the capital stock from the destruction of World War II, and were bound to end as the continental European and Japanese economies converged towards US levels of output per capita.

The UK and Canada display a more stable profile, similar to that of the US, although the slowdown in the 1970's is more clearly visible in the Canadian data. Again, similar to the US, an acceleration in the late

Figure 6. Posterior estimate of long-run growth (1960-2014)

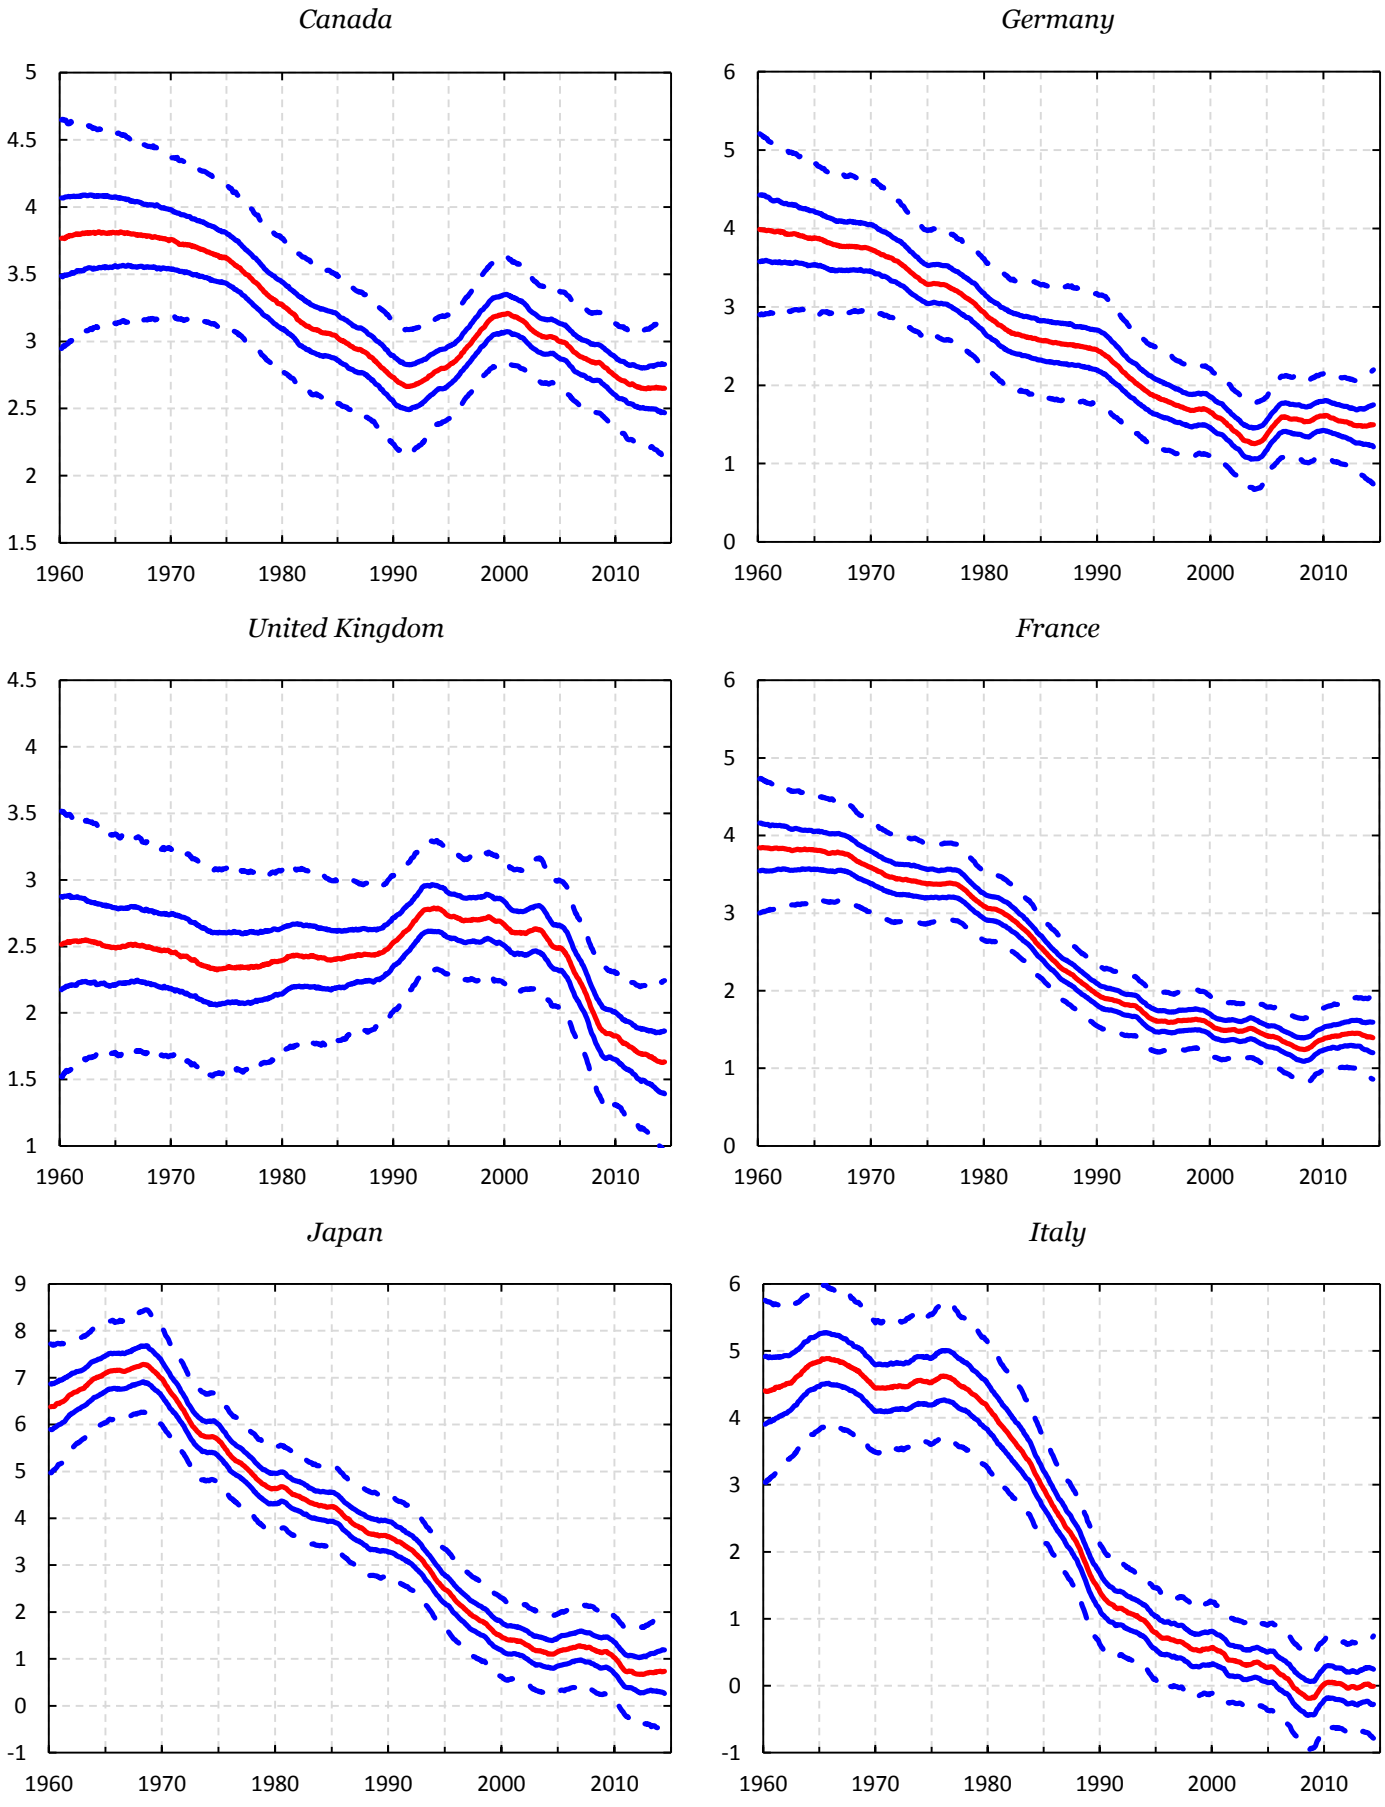

1990's and a subsequent slowdown in the mid 2000's is observed. It is interesting to note that applying the break test to the other G7 economies breaks are detected for most countries, and once again breaks are clustered around two episodes, the early 1970's and the early 2000's. An exception is the UK, where no break is detected but our model estimates a substantial decline in trend growth, an issue that has been extensively discussed in UK policy circles. Our real-time application to the US suggests that it is possible that our model is detecting the decline in long-run growth earlier than the break test.

By applying a simple accounting identity, we can take a first step at giving a more structural interpretation of the estimated fluctuations in long-run growth. By this identity, the growth rate of output is equal to the sum of the growth rates of labour productivity, hours per capita and population growth. This exercise reveals, as expected, that secular movements in the growth rate of productivity are behind the bulk of changes in long-run growth. Sometimes, however, other factors can play an important role. In Japan and continental Europe, population growth slowed down during the sample 1960-2014, making a negative contribution to overall growth in the last decade in Germany and Japan, and close to zero in France and Italy.

In the US, the "productivity slowdown" of the early 1970s was partially masked by an increase in hours per capita (mainly a result of increases in female labour force participation) and population growth during the period 1973-2005, resulting in a broadly stable long-run growth rate of about 3% that has come to be regarded as a stylized fact of the US economy. However, as pointed out by Gordon (2014), in recent years, the growth rates of productivity, hours per capita and population have all slowed down persistently, leading to a clear decline in long-run GDP growth significant enough to be detected by structural break tests.

It is worth noting that the weakness of productivity of recent years is not confined to the US economy. Figure 7 illustrates this point by plotting the five-year centered moving average of the growth in real output per hour worked. A marked slowdown is visible in all countries, and all of the areas are experiencing productivity growth below 1%, an unprecedented phenomenon in the post-War period. The case of the UK, where measured productivity has been slightly negative since the Financial Crisis, is particularly striking. The coincidence in the timing of the current productivity slowdown across countries suggests a driving common factor, and a more `structural' interpretation of the decline in productivity growth remains an interesting open question which we leave for further research. Nevertheless, the possibility that trends in productivity growth are substantially different currently than they were historically highlights the need to incorporate the uncertainty about long-run growth into econometric models and projections.

Figure 7. Labour productivity growth, % 5-year centered m.a.

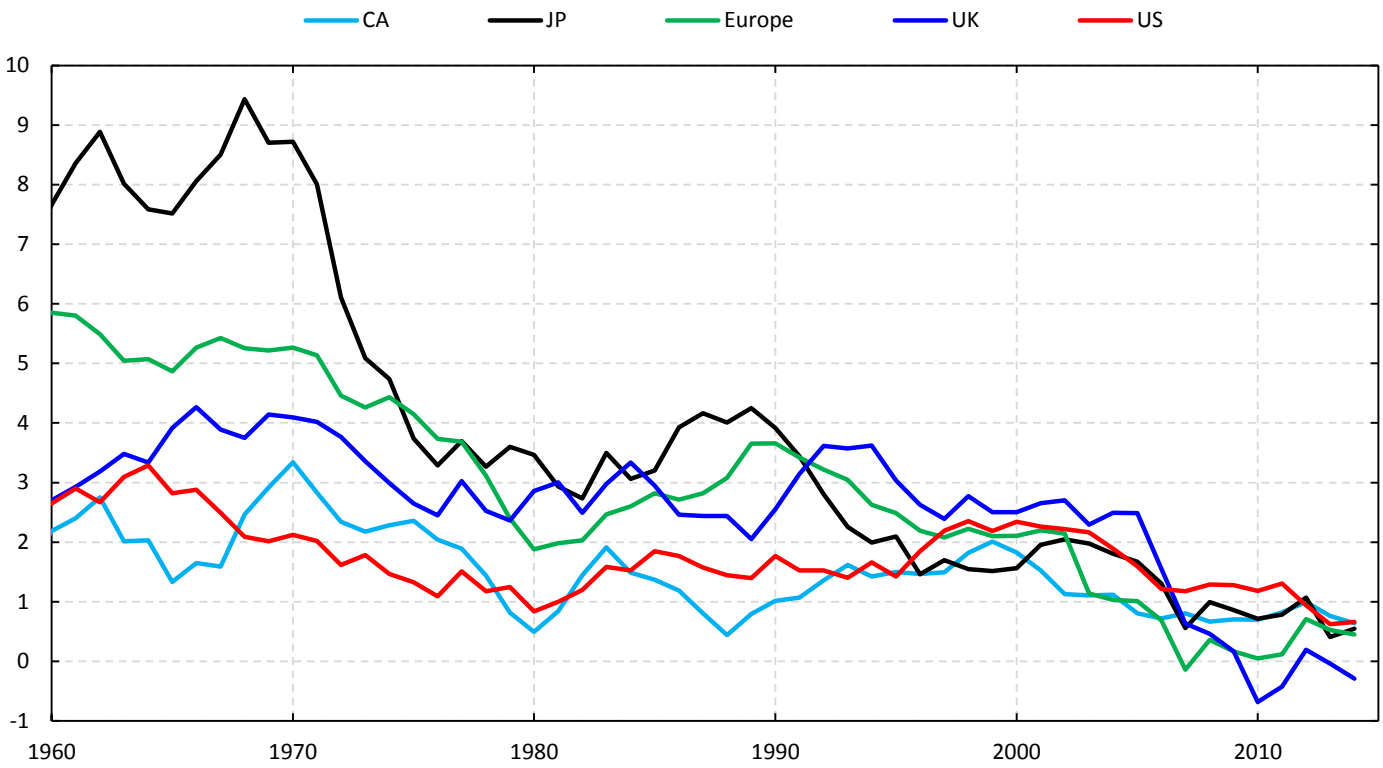

Note: Labour productivity calculated as the ratio of total real output to aggregate hours worked. "Europe" refers to the GDP weighted average of Germany, France and Italy. Sources: Conference Board Total Economy Database, IMF World Economic Outlook.

References

Banbura, M., Giannone, D., Modugno, M., and Reichlin, L. (2012). "Now-casting and the real-time data flow," Working Paper.

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

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.

Gordon, R. J. (2014). "The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reections," *NBER Working Papers 19895*, National Bureau of Economic Research, Inc.

Luo, S. and Startz, R. (2014). "Is it one break or on going permanent shocks that explains US Real GDP?" *Journal of Monetary Economics*, 66:155(163).

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