Brookings Papers
ON ECONOMIC ACTIVITY

BPEA Conference Drafts, March 23–24, 2017

The disappointing recovery of output after 2009

John G. Fernald, Federal Reserve Bank of San Francisco
Robert E. Hall, Stanford University
James H. Stock, Harvard University
Mark W. Watson, Princeton University

JOHN G. FERNALD
Federal Reserve Bank of San Francisco

ROBERT E. HALL
Stanford University

JAMES H. STOCK
Harvard University

MARK W. WATSON
Princeton University

The Disappointing Recovery of Output after 2009

March 10, 2017 draft

ABSTRACT: U.S. output has been expanding only slowly since the recession trough in 2009 even though unemployment has declined as fast as previous recoveries. We use a quantitative growth-accounting decomposition to explore explanations for the output shortfall, giving full treatment to cyclical effects that, given the depth of the recession, should have implied unusually fast growth. We find that the growth shortfall has almost entirely reflected two factors: TFP has grown slowly and labor force participation fell. Both factors reflect powerful adverse forces largely—if not entirely—unrelated to the financial crisis and the U.S. recession. Indeed, these forces fairly clearly were in play before the recession. The noncyclical forces we study resulted in a shortfall of capital formation that holds back output even today.

We thank Larry Ball, Robert Barro, Vasco Curdia, Lucretia Reichlin, Glenn Rudebusch, John Williams, and seminar and conference participants at the San Francisco Fed, the Boston Fed, and the New York Fed. We also thank John Coglianese and Neil Gerstein for excellent research assistance. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco, the Board of Governors of the Federal Reserve System, or the Business Cycle Dating Committee of the National Bureau of Economic Research.

Why has output grown so slowly in the post-2009 recovery, given the normal or better-than-normal recovery in the labor market? The U.S. unemployment rate has recovered at least as fast as in previous cyclical expansions—see Figure 1, left panel, where the dashed lines show changes in the unemployment rate following the troughs of recent recessions. In contrast, the right panel shows that the growth of output after 2009 has fallen far short. Output per person—the black line, in logs—fell sharply in the recession and has not reverted to any linear trend line extending its pre-recession trajectory.

The red line removes the effects of the deep recession in a simple way using Okun's Law, as described later in this paper. Because the economy had approximately returned to full employment by mid-2016, we have normalized the lines so that the red line intersects the black line at the end of the sample. The picture is striking: Cyclically adjusted output per person rose only slowly after 2007 and then plateaued.

We argue for taking this red line seriously as the counterfactual path of output in the absence of the recession. What appears to be a slow recovery of output is a reflection of something quite different: The U.S. economy suffered a deep recession superimposed on a sharply slowing trend.

To reach this conclusion, we first use Solow-style growth accounting to tease out the various components underlying the flattening of the red line. The answer is slow growth in total factor productivity growth (TFP), and falling labor force participation. The decline in participation was large enough that cyclically adjusted hours worked per person fell sharply. When put together, slowly rising TFP and falling participation imply flat cyclically adjusted output per person. Second, we examine TFP and participation in detail to understand whether their path has been influenced by the post-2007 experience of recession and slow recovery. Our answer is no. These factors reflect powerful adverse forces largely—if not entirely—*unrelated* to the financial crisis and recession.

The forces of declining productivity and shrinking labor force were in play before the recession. For example, Jorgenson, Ho, and Stiroh (2008) and Oliner, Sichel, and Stiroh (2007) noted that TFP growth had slowed by 2008 from its exceptional pace from the mid-1990s to the mid-2000s. And the Congressional Budget Office (2006) and Aaronson and others (2006) forecasted declines in participation as the baby boom retired and the surge of women into the labor force during the 1960s, 1970s, and 1980s plateaued.

Although many forecasters cut their forecasts for longer-term (cyclically adjusted) growth prior to the recession, the magnitude of the slowdown in actual output growth surprised forecasters over and over. Figure 2 shows the median forecast paths of the unemployment rate and of GDP from the Survey of Professional Forecasters, for forecasts made annually in 2010 through 2015, using data through the end of the previous year. These forecasts consistently exceeded actual growth. Early in the recovery, forecasts of the decline in the unemployment rate were borne out but, starting in 2013, they understated the improvement. These forecasts are representative of other real-time forecasts by the Congressional Budget Office, the Federal Open Market Committee (Lansing and Pyle, 2015), and the Council of Economic Advisers.

Some commentators have attributed the growth disappointments to weak investment and an absence of normal capital deepening in this recovery. In our view, the apparent absence of normal capital deepening largely reflects the adjustment of the capital stock to a slower underlying trend rate of output growth. Indeed, by mid-2016, when the economy had effectively recovered, the capital-output ratio was close to its pre-recession trend line.

Our account leaves little room for explanations of slow growth in which demand shortfalls have persistent effects. It does leave room for demand factors that delayed the recovery. Two quantitatively important factors are the unusually slow growth of federal government purchases during 2012 through 2014, which we associate in part with the sequester; and the delay in the usual rebound of state and local government purchases, which we associate with the aftermath of the housing market collapse and the financial crisis. Absent such delays, output growth would have been faster earlier: the black line in Figure 1 would have intersected the red line sooner. But, looking back over the entire recovery, we conclude that the seeds of the disappointing growth in output were sown prior to the recession in the form of a declining participation rate and slow TFP growth. Indeed, the scaling back of consumption and investment plans in response to the slowdown in TFP growth could induce its own recessionary pressures beyond those from the financial crisis alone. Blanchard, Lorenzoni, and L'Huillier (2017) show that these contributions could be large, especially with interest rates at the zero lower bound.

One loose end is that, under standard growth theory, the decline in TFP growth and participation should result in a rise in the capital-output ratio—slow growth reduces the volume of investment needed to keep capital on the growth path. By 2016 the cyclically-adjusted capital-output ratio had returned to its trend growth path, but it did not rise above that path as growth theory would suggest. One possibility is that further adjustments could lie ahead with additional capital deepening. Or, as Gutiérrez and Philippon (2016) suggest, other non-cyclical factors since around 2000 are pushing down the steady-state capital-output ratio.

To complement our growth-accounting decomposition, we use a three-part time-series representation of the variables we consider. These comprise a cyclical part, a long-run low-frequency trend part, and a remaining irregular part.

We use two approaches to identify the cyclical part. The goal is to adjust our analysis of the recovery to account for the depth of the 2007-2009 recession—deeper recessions tend to be followed by bigger recoveries. The first approach follows Okun (1962) by regressing the growth of a variable on the change in unemployment. For example, for output, we expect that growth will be unusually high in a recovery, which is needed to bring the unemployment rate down. The regression on changes in the unemployment rate gives an estimate of that cyclical part. We use a benchmark based on the three preceding recoveries to make comparisons between the post-crisis recovery and evolution of output and other variables. We make the comparison separately for the trend, cycle, and irregular parts.

Our approach based on the use of unemployment as a cyclical indicator concludes that the growth rate of business output per person in the recovery fell short of the normal described by the three prior recoveries by 1.8 percent per year, cumulating to a total shortfall over the recovery of 13.5 percent. The shortfall of TFP growth contributed nearly 1 percentage point per year, cumulating to 7 percentage points of output shortfall. The shortfall of participation accounted for 0.9 percentage points per year of the output shortfall, cumulating to 6.1 percentage points of the shortfall in output.

Our second approach asks a related question: Was the recovery disappointing relative to expectations in 2009, at the trough of economic activity? We use a dynamic factor model to condition on the state of the economy at the trough and simulate a forecast that would have been made at that time. The shortfall of actual performance relative to that forecast measures the unusual aspects of the recovery.

We find that the annual growth of real business output per person over the period from 2009 through 2016 was expected to be 2.0 percent per year at the beginning of the recovery—a figure well below the rate in earlier recoveries—but actual growth was even less, at 1.7 percent per year, so there was a shortfall of 0.3 percentage points. Our decomposition shows a shortfall of 0.5 percentage points in growth of TFP, 0.2 percentage points from capital shallowing, and a net positive contribution from labor input of 0.4 percentage points. A good part of that contribution is the absorption of unemployed workers back into employment. But the decline in the labor force contributed 0.4 percentage points to the shortfall in output per person.

The centrality of the decline in TFP, and in the growth rate of the labor force participation rate, leads us to examine them in greater detail.

Total factor productivity. In Section IV, we examine the decline in TFP growth and the extent to which the slowdown in labor productivity growth can be attributed to a slowdown in capital deepening—that is, to capital shallowing. Using time series methods that adjust for normal cyclical movements, we find that the slowdown in TFP growth occurred before the recession; using regime-shift detection methods, we estimate a break date in early 2006. Alternative Bayesian estimates, which do not assume a sharp break, place the date even earlier. The time-series estimates are in line with declines seen in annual sectoral productivity data. The timing matters: If, as the empirical evidence suggests, the slowdown in cyclically adjusted TFP growth occurred before the recession, the recession cannot be its cause. Moreover, after cyclical adjustment, weak investment and capital growth does not appear to have been an important independent contributor to weak output growth over this recovery, and actual investment during the recovery was almost exactly in line with our simulated forecast at the beginning of the recovery. Although capital formation has been below par, so has output growth, and by 2016, the capital/output ratio was in line with its long-term trend. Finally, the log-linearization of labor productivity expresses its growth as the sum of the growth of TFP and the growth of the capital/output ratio. Because the growth of the capital-output ratio has returned to normal, TFP dominates the movement of labor productivity.

We are therefore left with the conclusion that the mid-2000s slowdown in TFP growth played a key role in the slow growth of output during the recovery. We review a number of candidate explanations for the mid-2000s TFP slowdown and provide some new evidence against one, namely changes in regulations. We lean toward the hypothesis that the slowdown reflects at least a pause in the broad-based, transformative effects of information technology—the productivity boom that began in the mid-1990s ended in the mid-2000s.

Looking ahead, a key question is whether the slow growth of cyclically adjusted TFP since the mid-2000s is an unlucky period which will revert to the higher, IT-led TFP growth of the previous ten years, or alternatively the period from the mid-1990s through mid-2000s was the lucky period, and the economy must now adapt to the lower growth of output and wages implied by persistently low TFP growth.

The labor-force participation rate. In Section V, we turn to the decline in the labor-force participation rate that occurred from 2010 to 2016. Some of the decline in participation was probably a lingering consequence of the rapid rise in unemployment during the recession, participation continued to decline during the recovery in spite of a large steady decline in unemployment and the corresponding improvement in the availability of jobs. By the end of 2016, the unemployment rate was a percentage point below its long-run average, yet in 2016 the participation rate had fallen to 62.7 percent, three percentage points below its value at the trough. Although different methods for estimating the cyclical component of the participation rate provide different estimates of its cyclical decline early in the recovery, by 2016 that cyclical contribution was small.

It has been widely observed that the retirement of the baby boom is an important factor behind the decline of the participation rate. Less widely recognized is that there are other factors pushing the other way, notably the increasing level of education of the newly older workers. We put together these factors with an index that allows for shifting population shares in age, education, gender, and marital status, and find that these demographic effects account for 0.6 percentage points of the overall decline of 1.8 percentage points during the recovery. Changes in participation rates within detailed demographic groups account for the remaining 1.2 percentage points, or nearly two-thirds, of the decline since the cyclical trough.

There is no consensus about the sources of the persistent unexplained component of participation. We believe that it is not plausibly a consequence of the increase in unemployment in the 2007-2009 recession. The twin recessions of the early 1980s raised the unemployment rate by a comparable amount, and the recovery of the unemployment rate starting in 2009 was comparably fast and complete to that starting in 1982, but there was no comparable decline in participation relative to trend. Our review of the evidence supports the less optimistic view, that the non-demographic part of the decline represents a continuation of pre-existing trends that have a variety of sources that are likely to persist.

Timing of the recovery and demand considerations. During the period when unemployment remained above normal, concerns developed that the zero lower bound, the limited scope of fiscal policy, and other factors, might result in persistent deficient demand. We use the dynamic factor model to study the detailed components of expenditure to shed light on the sources of deficient demand during the recovery. As in our earlier analysis, we calculate a simulated forecast as of 2009 and study its errors in terms of growth rates from 2010 to 2016. The errors are stated as percentage-point contributions to an overall forecast error of 0.57 percent of GDP per year, close to the Okun's law shortfall of 0.73 that remains after adjusting for slower trend growth and normal cyclical movements.

More than half of the total forecast error—0.31 percentage points per year—arises from shortfalls in government purchases of goods and services (0.20 federal and 0.12 state-local). Direct fiscal policy—government infrastructure purchases and the like—was a substantial factor restraining the expansion, relative to past experience as summarized in the forecast. Indeed, according to the factor model, government consumption expenditures plus transfer payments would normally have grown by 2.9 percent per year over this period, but in fact grew by only 0.7 percent per year, a shortfall of 2.2 percentage points. Examination of the forecast paths points to slow growth of state and local purchases in the first four years of the recovery, and weak growth of federal purchases proximate to the onset of the sequester.

Total household consumption—by far the largest component of total spending—contributed 0.26 percentage points per year to the shortfall in output growth. Durable goods, the most cyclical part of consumption, behaved almost exactly as forecast during the expansion, as did nondurable goods. Roughly half of the shortfall arose in two parts of services: housing and financial services. This finding supports the conclusions of a large body of research that has focused on housing and finance as key sectors for understanding the special features of the recession and recovery. In contrast, the real value of financial services is a particularly poorly measured component of output, and the shortfall in this sector, plus that in the even more-poorly measured sector of nonprofit institutions serving households, contributes fully 0.10 percentage points to the 0.57 percentage point under-forecast of output.

These forecasts suggest little role for some of the weak-demand explanations. The absence of any significant shortfall in consumption growth outside housing, is evidence against the hypothesis that deleveraging and increasing inequality contributed to the slow recover. Weak exports exerted a small drag on output growth, mainly during 2011-2013. And business investment was slightly stronger than the forecast based on earlier recoveries. This last finding supports our general hypothesis that business investment, a highly cyclical endogenous variable, behaved essentially normally in the recovery and is not an exogenous contributor to the weakness of the recovery.

I Growth Decomposition and Data

Section I.A describes our general objective and our data. Section I.B then lays out the Solow-style growth-accounting framework we use to analyze the slow recovery in output.

I.A Focus and Data

We focus on understanding the disappointingly slow recovery that started in mid-2009, when the National Bureau of Economic Research dates the end of the recession. We end seven years later, in 2016. When we make comparisons to the preceding three recoveries, we use the comparable seven-year periods following the troughs, except following 2001, when we truncate at the business-cycle peak at the end of 2007 (six years).

The slow recovery in output can be examined through the lens of production (output is produced) or expenditure (output is purchased). Here we discuss growth-accounting identities related to production. The production framework is natural for addressing the role of structural trends such as productivity and the labor force. We apply this accounting to the business sector. Growth accounting is less applicable to government, household, and non-profit production, where output is often not measured independently of inputs.

Our measure of output is the geometric average of income and expenditure side measures, as recommended by the recent literature—see the data appendix. Both sides of the accounts provide information about true growth but are subject to measurement error, so a combination improves the signal-to-noise ratio. At an economy-wide level, we refer to this average of gross domestic product and gross domestic income as gross domestic output (GDO) or, where the context is clear, just output. Unless noted otherwise, we scale output by the population eligible for employment, aged 16 and above, denoted *Pop.*

Our business-sector growth-accounting data are described in Fernald (2014). These quarterly data provide the values of the variables in the equations below. Broader real gross product and gross income aggregates come from the Bureau of Economic Analysis, and labor-market data from the Bureau of Labor Statistics. The data appendix provides further details.

I.B Accounting for Growth

Although our growth accounting focuses on the business sector, we need to consider the overall economy because labor market indicators, such as the unemployment rate, measure that concept. Identities link economy-wide gross domestic output, *GDO,* and business output, Y_t^{Bus} :

$$\left(\frac{GDO_t}{Pop_t}\right) = \left(\frac{GDO_t}{Y_t^{Bus}}\right) \times \left(\frac{Y_t^{Bus}}{Pop_t}\right) \tag{1}$$

The identities in this section are sometimes in levels, sometimes in growth rates, depending on which is clearer. Empirical estimation is in growth rates.

Growth accounting decomposes output growth into a set of components that help to show how the second term in equation (1) evolves. Modern growth accounting follows Jorgenson and Griliches (1967) which, in turn, expanded and clarified Solow (1957). Growth in business output, Y_t^{Bus} , depends on growth in capital, *K,* and labor input, *Labor*. Labor, in turn, depends on *Hours* and labor quality, *LQ:* $\Delta \log Labor_t^{Bus} = \Delta \log LQ_t + \Delta \log Hours_t^{Bus}$. Labor quality *LQ* captures the contribution of rising education and experience. Our measure of *LQ* assumes that relative wages capture relative productivities of workers with different attributes—see Bosler and others (2016). In per-person terms, we write:

$$\Delta \log\left(\frac{Y_t^{Bus}}{Pop_t}\right) = \Delta \log TFP_t + \alpha_t \Delta \log\left(\frac{K_t}{Pop_t}\right) + (1-\alpha_t)\Delta \log\left(\frac{LQ_t \cdot Hours_t^{Bus}}{Pop_t}\right) \tag{2}$$

The time series α_t is capital's share of income.

For some purposes, we rewrite equation (2) in a way that distinguishes endogenous from exogenous factors. For example, suppose a demographic change reduces growth of hours of work. In equation (2), that change is multiplied by labor's share. But if the same force that cut hours of work also affected capital input, as growth models generally predict, we may want to incorporate the endogeneity of capital. For this purpose, we consider an alternative decomposition of (Y_t^{Bus} / Pop_t) as business-sector hours per person times labor productivity (output per hour of work):

$$\left(\frac{Y_t^{Bus}}{Pop_t}\right) = \left(\frac{Hours_t^{Bus}}{Pop_t}\right)\left(\frac{Y_t^{Bus}}{Hours_t^{Bus}}\right) \tag{3}$$

The first term on the right-hand side, business hours per person can be expanded as:

$$\left(\frac{Hours_t^{Bus}}{Pop_t}\right) = \left(\frac{Hours_t^{Bus}}{Emp_t^{Bus}}\right) \times \left(\frac{Emp_t^{Bus}}{Emp_t^{CPS}}\right) \times \left(\frac{Emp_t^{CPS}}{LabForce_t}\right) \times \left(\frac{LabForce_t}{Pop_t}\right) \tag{4}$$

The terms on the right-hand side of (4) are as follows:

- $\left(\frac{Hours_t^{Bus}}{Emp_t^{Bus}}\right)$ is business-sector hours per employee.
- $\left(\frac{Emp_t^{Bus}}{Emp_t^{CPS}}\right)$ is the ratio of business employment, measured (primarily) from the establishment survey, to household employment, measured from the Current Population Survey (the household survey).[1]

[1] There are several conceptual differences between business employment and household employment, in addition to the source data. A quantitatively important one is that the household survey covers the entire civilian economy, which is broader than the business sector, and, correspondingly, is less cyclical. Fernald and Wang (2016) discuss differences between the business-sector and household-survey measures and why this gap is procyclical. (so it tends to fall when the unemployment rate rises). They find that, once the coverage differences are taken into account, the cyclicality of total hours worked is similar between the two surveys.

- $\left(\frac{Emp_t^{CPS}}{LabForce_t}\right)$ is employment relative to the labor force, and is by definition equal to 1-U_t, where U_t is the unemployment rate. Over the long run the contribution of the U term is zero because the unemployment rate reverts to a mean value between 5 and 6 percent.
- $\left(\frac{LabForce_t}{Pop_t}\right)$, the final term, is the labor force participation rate.

Now consider labor productivity, the second term on the right-hand side of equation (3). With some manipulation, the growth-accounting equation (2) yields a useful expression:

$$\Delta\log\left(\frac{Y_t^{Bus}}{Hours_t^{Bus}}\right)=\frac{\Delta\log TFP}{(1-\alpha_t)}+\left(\frac{\alpha_t}{1-\alpha_t}\right)\cdot\Delta\log\left(\frac{K_t}{Y_t^{Bus}}\right)+\Delta\log LQ_t\,. \tag{5}$$

In this expression, output per hour depends on the capital-output ratio, and labor quality, both expressed in labor-augmenting form. It is useful because we tend to interpret capital deepening as endogenous. With slower growth in technology and labor, the path of capital will be lower—the capital/output ratio will remain roughly stable. Thus, the ratio is useful in assessing whether there is a special influence on capital, for example from unusual credit constraints or from heightened uncertainty. The baseline is not the level of capital but the capital/output ratio.

In the one-sector neoclassical growth model, the capital-output ratio is pinned down by an Euler equation. If trend technology were constant, the steady-state ratio is stationary. In models with investment-specific technical change—and in the data—that ratio has a relatively slow-moving trend—see the online appendix to Fernald (2015).

Of course, the capital-output ratio is not necessarily dispositive. A reduction in trend technology raises the steady-state capital/output ratio, which then pushes down the equilibrium real interest rate. Other factors, such as an increase in market power (e.g., Gutiérrez and Philippon, 2016) could work in the other direction. Nevertheless, in the data, the trend capital/output ratio estimated from cyclically adjusted data has been remarkably smooth since the 1970s, despite the speedup in growth in the mid-1990s and the slowdown in the mid-2000s. We conclude that the capital/output ratio is informative about the possibility of a capital shortfall.

II Estimation of Cyclical Components and Low-Frequency Trends

After the economy reaches a cyclical trough following a negative shock, the rate of unemployment returns to a normal or natural rate, and while this is happening, output grows faster than it would with constant unemployment. The larger the negative shock, the greater the recovery in the labor market and the greater the cumulative above-normal growth of output. Thus, in determining whether the recovery from the 2007-2009 recession was slow, we need to control for the depth of the recession. Moreover, the calculation needs to control for underlying systematic changes in the U.S. economy, such as changes in immigration and the demographics of the workforce, that affect the underlying mean growth rate of employment and output.

In this paper, we use two complementary methods for controlling for the depth of the 2007-2009 recession and thus for assessing the speed of the recovery. The first method conditions on the path of unemployment. This method asks the question, what would the normal cyclical path of output and the other variables in the growth decomposition have been, given the 2009-2016 recovery in the unemployment rate? In practice, this amounts to estimating the normal cyclical movements using Okun's Law, extended to variables in addition to output.

The second method controls for the depth of the recession by conditioning on the state of the economy at the 2009 trough, as measured by a large number of time series,. This method asks the question: What would the normal cyclical path of output, the growth decomposition variables, and other macroeconomic variables have been, given the depth of the recession in 2009? Calculating the normal path involves simulating forecasts of multiple time series, given data through 2009, and for this purpose we use a high-dimensional dynamic factor model.

Both methods allow for low-frequency changes in mean growth rates, that is, for trends in the growth rates. To this end, throughout this paper, we adopt a statistical decomposition of the growth rate of a given time series into a trend, cycle, and irregular part. Let y_t be the percentage growth rate of a variable at an annual rate, computed using logs. For example, for GDO, $y_t = 400\ \Delta\log GDO_t$. The statistical decomposition is

$$y_t = \mu_t + c_t + z_t, \tag{6}$$

where μ_t is a long-term trend, c_t is a cyclical part, and z_t is called the irregular part—it describes the higher-frequency movements of the variable that are not correlated with the cycle.

Following convention in the time series literature, we refer to equation (6) as a trend-cycle-irregular decomposition. Because y_t is a growth rate, the trend μ_t is the long-term mean growth rate of the series. In the special case that this mean is constant, in log-levels the series would have a linear time trend, with a shifting intercept that depends on c_t and z_t. As explained below, we estimate the long-term trend as the long-run average of y, after subtracting the cyclical part. This long-run average typically changes over time—for reasons such as changing demographics. Our quantification finds that those changes are important for understanding the weak recovery in output.

The irregular term, z_t, is the variation in y_t net of the trend and cyclical fluctuations. In the context of this paper, this irregular term is of central interest: It represents the shortfall or excess of the growth in a given variable during the recovery, above and beyond what would be expected given low-frequency changes in the economy such as demographics, and the normal cyclical movements expected during the recovery from a deep recession. We find large negative irregular parts play important roles in the weak recovery.

II.A Method 1: Using Okun's Law to Account for the Cycle

The first method uses Okun's Law to extract the cyclical component. Because we consider many series, and those series often lead or lag the unemployment rate, we extend Okun's relationship to include leads and lags. The Okun's Law definition of c_t thus is,

$$c_t = \sum_{j=-p}^{q} \beta_j \Delta u_{t+j} = \beta(L)\Delta u_t \tag{7}$$

where u_t is the unemployment rate and β(L) is the distributed lag polynomial with q leads and p lags in the summation. Choice of p and q and other estimation details are discussed in the next subsection. The sum of the lag coefficients, $\beta(1)$, is a measure of the overall cyclical variability of y_t. Note that because $E\Delta u_t = 0$ over the long run, our cyclical part has long-run mean zero.

Okun's original relationship was the reverse regression of changes in the unemployment rate on changes in output with only contemporaneous movements. However, subsequent researchers have often used the specification with unemployment on the right-hand side, and for output growth and many other series, the leads or lags or both are statistically significant, so we refer to equation (7) as a generalization of Okun's law.

The standard unemployment rate is only one of many measures of the state of the labor market. Other plausible indicators include marginally attached workers, workers working part-time for economic reasons, discouraged workers, the long-term unemployment rate, and the short-term unemployment rate. One can imagine adding such measures to equation (7). However, using the standard unemployment rate, as we do here, has several virtues. It is well-measured, and has been measured using essentially the same survey instrument since 1948. Over the long run, it has essentially no trend. And in any event the other measures of the state of the labor force are highly correlated with the unemployment rate, once one incorporates leads and lags. For example, Figure 3 shows one alternative measure, part-time workers as a fraction of employment, a series which moves closely with the unemployment rate.

An alternative approach to measuring the cycle would be to condition on the path of the output gap instead of the unemployment rate. But conditioning on output would prevent addressing the question of this paper, why output growth has been slow during this recovery. Rather, our paper seeks to understand why output growth has been so weak after taking into account the depth of the recession and the strength of the labor market recovery.

Cyclically adjusted trend. A practical problem in estimating the trend μ_t is that persistent cyclical swings can be confused with lower frequency trends. This problem is particularly acute in estimating trend terms towards the end of our sample given the severity of the recession and length of the recovery. To address that problem, our estimate of the trend controls for normal cyclical movements implied by Okun's Law. That is, we do not estimate the trend directly from the variables, but rather after removing the cyclical component from our earlier regression.

Substitution of equation (7) into equation (6) yields

$$y_t = \mu_t + \beta(L)\Delta u_t + z_t\,. \tag{8}$$

The Okun's Law "residual" (including μ_t), $y_t - c_t = y_t - \beta(L)\Delta u_t$, is a measure of what growth rate would have been consistent with an unchanged unemployment rate. To estimate μ_t, we adopt the framework of the partially linear regression model, which treats μ_t as a nonrandom smooth function of t/T; see Robinson (1988), Stock (1989) and Zhang and Wu (2012). In this approach, μ is estimated as a long-run smoothed value of y, after subtracting the estimated cyclical part:

$$\hat{\mu}_t = \kappa(L)\left(y_t - \hat{\beta}(L)\Delta u_t\right) \tag{9}$$

where κ(L) is a filter that passes lower frequencies and attenuates higher frequencies. Because the estimated cyclical part is subtracted prior to smoothing, we will refer to the estimated trend $\hat{\mu}_t$ as a cyclically-adjusted trend. The use of a cyclically adjusted trend with a long bandwidth for κ(L) helps avoid attributing the recent slow growth mechanically to a declining trend. The Econometric Appendix compares the partially linear regression approach to a state space (or unobserved components) methods, and discusses computation of the heteroskedasticity- and autocorrelation-robust standard errors.

Estimation. We estimate β(L) by the least squares regression of y_t on leads and lags of Δu_t, where u_t is the unemployment rate. We chose $p = q = 2$ based on sensitivity analysis: For some left-hand variables, using only contemporaneous Δu_t suffices, but for others additional leads and lags are justified statistically. Our overall results are robust to using more leads and lag. Our estimation period starts at the 1981 peak and ends in mid-2016.

For the low-pass filter κ(L), we use a biweight filter with truncation parameter of 60 quarters. Tukey's biweight filter $\kappa(L)$ is two-sided with $\kappa_j = c(1—(j/B)^2)^2$ for $|j| \leq B$ and = 0 otherwise, where B is the bandwidth and c is a normalization constant such that $w(1) = 1$. End points are handled by truncating the filter outside the range of the data and renormalizing. The long truncation parameter—the filter weights span ±15 years—was chosen so that changes in $\hat{\mu}_t$ reflect slow multi-decadal swings. If there are sharp shifts or breaks in trend growth, this filter will over-smooth, an issue we consider in discussing the evolution of TFP.

Additivity. The foregoing method for estimating the trend, cycle, and irregular parts has the useful property that it preserves additivity when applied to additive decompositions. Specifically, suppose that $y_t = y_{1t} + y_{2t}$. This additivity is preserved for the estimated cyclical, trend, and irregular parts: $\hat{\mu}_t = \hat{\mu}_{1t} + \hat{\mu}_{2t}$ and $\hat{c}_t = \hat{c}_{1t} + \hat{c}_{2t}$, where the subscripts refer to the parts of y_t, y_{1t}, and y_{2t}. This property is a consequence of using the same cyclical regressors and same filter $\kappa(\mathrm{L})$ for all series, and the property that regression is linear in the dependent variable.

II.B Method 2: Dynamic Factor Model Estimates of the Cycle

The dynamic factor model produces forecasts of the variables under study using the history of a broad cross section of macro variables through the trough in 2009, second quarter. A small number of common factors are extracted from 123 macro variables, and these are used to summarize the state of the macro economy in 2009. Forecasts of the factors show how the state of the economy would have been predicted to evolve based on the history of the factors. The factor forecasts are then used to forecast the series of interest using the historical correlation between the series and factors.

Stock and Watson (2016) discuss factor methods and provide extensive results for an empirical factor model using a closely related large dataset. Here, we briefly summarize key steps—for additional detail, including variable transformations and measures of fit, see the online appendix to this paper.

Estimation of the factors and dynamic factor-model parameters, and computation of forecasts. We work with the static form of the factor model,

$$X_t = \Lambda F_t + e_t \tag{9}$$

$$F_t = \Phi(L)F_t + \eta_t, \tag{10}$$

where Λ is the $N \times r$ matrix of factor loadings, $\Phi(L)$ is the vector autoregression lag polynomial for the factors, X_t is the $N \times 1$ vector of series, F_t is the $r \times 1$ vector of factors, and η_t are the innovations to the factors. The term ΛF_t is referred to as the common component of X_t and e_t is the idiosyncratic component.

The series used to estimate the factors are summarized in Table 1, and the full list is provided in the data appendix. The dataset omits high-level aggregates to avoid aggregation identities and double-counting—for example GDP is omitted, because its components are included, consumption of goods is omitted because durables and nondurables consumption are included separately, and total employment is omitted because its components are included.

The 123 series are used to estimate six factors by principal components. The factor loadings Λ_i of series i are then estimated by regressing X_{it} on the estimated factors using data from 1984 through the trough. The six factors themselves are forecasted using a vector autoregression with 4 lags, with a jumping off point of the trough quarter, producing a series of forecasts of the factors $\hat{F}_{t|2009Q2}$ for successive quarters through mid-2016. With those factors in hand, forecasts for variable i are computed as $\hat{X}_{t|2009Q2} = \hat{\Lambda}_i \hat{F}_{t|2009Q2}$ for those quarters. The details are given in the Econometric Appendix.

Post-trough trends for forecasts. The simulated forecast approach involves freezing the trends in each series at their trough values (with one exception), and projecting a constant trend growth. The exception is that we allow for demographic changes to affect labor force participation. It was recognized before the recession that the imminent retirement of the baby boom would depress participation, see Aaronson and others (2006) and Congressional Budget Office (2007). Here, we use a Divisia-Tørnqvist index to project the effect of evolving demographics, specifically the effect on overall participation of changes in the population shares by age, education, and gender. This index improves on the index that is common in this literature, which allows for changes only in age shares and uses age-specific participation rates for some fixed base year such as 2007. We defer further discussion of this index to Section V.

The projected demographic trend in participation feeds through, with share weights as appropriate, into the trends in employment, hours, and output. We leave the trends in capital, the ratio of business to household employment, and hours per employee unchanged. The result is an output trend that incorporates aging and other demographic effects on employment as understood at the trough, with other component trend growth rates frozen at their trough values. Trend growth rates of the demand components of output are computed as the component's time series trend as of the trough, plus the share-weighted difference between the output trend (inclusive of the participation aging trend) and the trough value of the output trend. This final adjustment ensures that the share-weighted trend growth rates add, however it is numerically negligible because the trough-quarter participation adjustment to the trend value of output is small.

***Additivity*.** The DFM, like the Okun method, preserves additivity of components.

III Results: Accounting for Slow Growth

We are now ready to quantify the sources, in a growth-accounting sense, of the slow growth in output. We begin with a brief discussion of the cyclical properties of the component variables in the growth-accounting decomposition.

III.A Cyclical Properties of the Growth-Decomposition Variables

Table 2 provides three summary measures of the cyclicality of the variables entering the growth decomposition and additional broad measures of output. The first is the generalized Okun's Law coefficient, the sum of the coefficients, $\beta(1)$, in equation (8), which is divided by four to yield standard units of percent change in output per percentage point change in the unemployment rate. Like the other parts of the decomposition, the generalized Okun's law coefficients also satisfy additivity, so that the sum of the Okun's law coefficients on the components equals the Okun's law coefficient on the sum of the components; that is, the coefficients in lines 7-9 add to -2.02, the coefficient on real business output per capita. The Okun's law coefficients in the first column provide a natural measure of the cyclical variation in business output per capita.

Of the total cyclical variation of business hours per person (line 10), as measured by the generalized Okun's law coefficient of -2.3, nearly half (-1.08) comes from the employment rate (one minus the unemployment rate), one-sixth (-0.35) comes from variations in hours per worker, and a small amount (-0.16) comes from labor-force participation. These results support the view that participation is slightly procyclical, falling as unemployment rises. Of course, a large unexpected reduction in participation occurred before and during the recovery. Section V asks whether the recent decline in participation can be related to the slack labor market.

One-third of the cyclical variation in business output (-0.71) comes from cyclical variation in the ratio of business employment to household-survey employment. When unemployment rises, business employment falls relative to economy-wide employment, as measured by households. Some of this difference arises from the higher stability of the non-business sectors. And some may arise from a cyclical discrepancy between the employment counts obtained by surveying business and non-business employers and counts from the CPS, for example a worker holding two jobs counts twice in the establishment survey but just once in the household survey. Fernald and Wang (2016) find that hours worked has almost the same cyclicality in the two surveys.

Labor productivity, line 15, is weakly and insignificantly countercyclical over our sample. It combines TFP (line 7 or, rescaled, line 16), which is strongly procyclical, with the capital-output ratio (line 17), which is strongly countercyclical. Research on TFP has discussed the roles of labor hoarding, cyclical changes in capital utilization, measurement errors, and other non-technological factors that account for the pro-cyclicality of productivity (see Basu and Fernald, 2001). Investment is pro-cyclical, but the cumulated stock of capital changes relatively little in synchrony with unemployment, so the capital-output ratio is strongly countercyclical because of output in the denominator. Finally, the countercyclicality of labor quality (0.13, row 18) supports the hypothesis that times of high unemployment are times of higher labor quality, because lower skilled workers differentially become unemployed.

The remaining columns of Table 2 quantify the amount of variation in the variable that is cyclical, as measured by first by the standard deviation of the Okun's law estimate of c_t and second by the fraction of the variance of the series explained by the factors (that is, the R^2 of the common component in the dynamic factor model). By both measures, the most cyclical variable is the employment rate—by construction for the Okun's law estimate and as a result of the factors explaining variation in employment for the factor-model estimate. Although cyclical variation in TFP accounts for one-fourth of the cyclical variation in business output per capita, cyclical variation only accounts for a fraction of the variation in TFP growth. TFP growth has a large amount of high-frequency variation, including measurement noise.

Figure 4, Figure 5, and Figure 6 show (in black) the log levels of the series in Table 2. These figures also plot the cyclically adjusted series, using Okun's Law (red), and the cyclically adjusted trend (blue). The black and red lines in the right panel of Figure 1 and in Figure 4(a) are the same, but with different time scales and normalizations (as in the figure notes).

III.B Growth Components: Trend and Cyclical Parts

Table 3 summarizes the results of the growth accounting decomposition, where Okun's Law is used to estimate the cyclical component conditional on the unemployment rate path. The table compares the mean values of these components in the recent recovery to their mean values in the three previous recoveries. For this table, the three previous recoveries are defined as the first 28 quarters of the recovery (the number of quarters from the first one after the trough to the end of our sample) or the trough-to-peak period, whichever is shorter. The left panel of three columns in the table presents actual average historical growth rates, and contributions to growth rates, at annual rates. The right panel, the remaining four columns, provides the decomposition after cyclically adjusting these variables using the Okun's-law method.

Table 4 is the counterpart of Table 3, in which the cyclical component is computed using the factor-based method, conditional on the state of the economy in mid-2009. The first column, the forecast, is the sum of the cyclical component of the forecast and the trend, averaged over the 2009-2016 forecast period. The second column is the actual average growth of the variable, and the third column is the factor estimate of the irregular part *z*, which is the shortfall, that is, the gap between the forecast and the actual. The standard error of the cyclical component (that is, the standard error of $\hat{\Lambda}_i \bar{\hat{F}}_{t|2009Q2}$ in the notation of equation (9)) is given in the final column.[2]

Figure 7 shows the forecasted and actual paths of the growth-accounting decomposition in Table 4, where the forecasted paths are computed using the factor model. The gap between the predicted and actual is the irregular part z_t, which is the forecast error from the factor model given the state of the economy at the trough.

III.C NIPA Expenditure Components: Trend and Cyclical Parts

Many proposed explanations for the slow recovery appeal to deficient demand, or in some component of demand. To shed light on these explanations, we therefore undertake an additional decomposition, this time based on the National Income and Product Accounts GDP expenditure identity stated in terms of its trend, cyclical, and irregular parts. The methods applied to the growth-accounting identities apply directly to the expenditure-account identities, and preserve additivity and internal consistency (up to log-linearization approximation).

Table 5 presents a decomposition of the forecast of output and its main product components. The entries in this table are contributions to mean growth, computed using share weighting; the entries in the first column correspond directly to Table 2 in the Bureau of Economic Analysis's release, *Contributions to Percent Change in Real Gross Domestic Product*,

[2] The shortfall in the third column is the negative of the usual definition of a forecast error. In addition, the standard error of the conditional mean in the fourth column is not the forecast standard error (which incorporates uncertainty associated with future values of the factors and shocks), but rather a summary of the sampling error associated with the estimated vector autoregression and other regression coefficients.

except that here the contributions are averaged over the 2009-2016 forecast period. Because the forecasts and forecast errors are additive, the trend values, forecasts and forecast errors in the remaining columns also add to their respective aggregates. The second block of columns in Table 5 presents results using the Okun's Law method for cyclical adjustment, and the right-hand block presents results using the dynamic factor model, so that the shortfall is the negative of the forecast error. Figure 8 and Figure 9 present additional plots, in the format of Figure 7, of the forecasted and actual values of selected variables in Table 5 and for employment growth, computed using the factor model. These series are not share-weighted.

III.D Discussion

A key difference between our two methods concerns the counterfactual cyclical path of labor market variables. Because the first method conditions on the unemployment rate path, by construction there is no irregular part for the unemployment rate, and the irregular part for closely related variable such as establishment employment is small. In contrast, the cyclical path in the forecasting exercise projects a normal cyclical path for all the variables, conditional on the state of the economy at the trough in 2009, and in principle, the actual path of any variable, including labor market variables, can depart arbitrarily from its forecast path. We find that the factor forecasts under-predict the robust recovery in the labor market and over-predict the growth of output. This recovery of employment combined with the slow growth in output is a key feature of this recovery. We return to this and other implications of the factor forecasts in more detail at the end of this next section.

Aside from this major difference in forecasts of output and unemployment in the recovery, the two methods generally yield quantitatively similar estimates of the irregular part, and lead to similar conclusions about the behavior of the components of output growth over the recovery. For clarity, we therefore focus primarily on results using the Okun method.

We begin with the first block of columns in Table 3, which summarizes the shortfall of output and the growth decomposition components without cyclical adjustment. GDO grew 3.57

percent per year in the previous three recoveries (column a), but only 2.20 percent in the current recovery (column b), for a shortfall of 1.37 percentage points (column c). Similarly, business output per capita grew 2.92 percent in the previous three recoveries, but only 1.72 percent per year in the current one, for a shortfall of 1.21 percentage points. Looking down column c, many of the rows are non-zero but a few stand out. These include a decline in the growth of capital per person (capital shallowing, row 8), a decline in the growth rate in TFP (rows 7 and 16), and a decline in the participation rate (row 14).

This comparison of actual growth rates understates the output shortfall, however, since it does not account for how deep the recent recession was relative to the three previous ones on average. The second block of estimates presents the same decomposition after removing the cyclical component using Okun's Law, that is, conditioning on the unemployment rate.

Making this cyclical adjustment creates a different, starker picture of the slow growth. The shortfall in business output per person is much larger, at 1.81 percentage points, reflecting the depth of the 2007-9 recession. The cumulative shortfall in output over this recovery is 13.5 percent (final column). Okun's Law cleans out the cyclical differences in many variables. The only element that is quantitatively important for explaining hours is labor-force participation (row 14). The only element that is quantitatively important for labor productivity is TFP (row 16). Shortfalls in the direct contribution of capital input per person are also large (row 8), but when scaled by output (row 17) the contribution is small.

We now discuss selected elements of the accounting.

Business output. Figure 4, which shows the cumulative parts of the growth of business output per capital, conveys a basic finding of this paper. For the period of the recovery from the crisis recession, the strong growth in the labor market should have been associated with a dramatic recovery in output, based on historical cyclical patterns. Indeed, as can be seen in Figure 1, the recovery in unemployment was essentially as rapid and complete as previous recoveries. But two powerful forces opposed the cyclical part—the low-frequency trend and the high-frequency irregular part. Moreover, the downward slopes of the two parts are almost the same, and our breakdown of the non-cyclical behavior of output gives equal roles to the high- and low-frequency parts.

Hours per worker. Figure 5 shows the levels of the three statistical parts of weekly hours per worker. Consistent with the coefficient of -0.35 in Table 2, the cyclical part of hours rose smoothly during the recovery, as in the three earlier recoveries. The slope of the low-frequency trend plotted in the figure, μ_t, rose slightly, while the high-frequency irregular part fell slightly. Unlike many other indicators, weekly hours behaved fairly normally in the post-crisis recession.

Labor force participation. The labor-force participation rate is the ratio of the sum of employment and unemployment to the population 16 and over. Figure 5 shows that the low-frequency trend in participation grew at a declining rate until 1998 and began to shrink after that year. The rate of shrinkage declined slightly in the last years shown. The cyclical part grew during the recovery, reflecting the small procyclical coefficient in Table 2, but both the high- and low-frequency parts declined. The net effect was a substantial decline in participation during the recovery, in contrast to the typical low but positive growth in recoveries. Section Vl pursues explanations of the anomalous behavior of the labor force during the post-crisis recovery.

Finally, Figure 6 shows that labor quality was an important part of non-cyclical movements of through the late 1970s and a positive contribution until the mid-1990s. Although the growth of labor quality slowed during the recent recovery, after cyclical adjustment this slowdown makes only a small contribution, 0.06 pp, to the slow growth of output.

Capital input. Capital input (row 8) contributes a moderate amount of non-cyclical movement to output. Its cyclical contribution is essentially zero. A decline occurred in the low-frequency part starting somewhat before the crisis. A small decline in the high-frequency part occurred during the same period.

As we noted earlier, capital input is jointly determined with TFP, the labor force, employment, and other endogenous variables. The movements of capital input per person shown in the figure reflect the joint determination of the variables. If the economy is hit by an exogenous decline in the rate of growth of TFP, optimal capital input grows less rapidly as well, according to most models of investment. Row 17, column h, shows that when stated relative to output, the shortfall of capital per unit of output after the crisis disappears.

Comparison of Okun's law and factor model shortfall estimates. Table 4 shows that compared to what would have been expected based on the data through 2009, actual GDP growth fell short by 0.57 percentage points, GDO growth by 0.43 percentage points, and business output by 0.35 percentage points. These cyclically-adjusted shortfalls are smaller than their counterparts in Table 3 because the recovery in employment was stronger than expected based on the factor forecasts. Whereas the Okun's law method in Table 3 conditions on the unemployment rate path, the factor model forecast has a shortfall in the CPS employment rate of -0.42 (the factor model predicts a less rapid fall in the unemployment rate). This feature of the factor forecasts—an unexpectedly strong recovery in the labor market and an unexpectedly weak recovery in output—is consistent with the forecast errors made in real time by professional forecaster evident in Figure 2. As a back-of-the-envelope comparison, using the Okun's law coefficient of 2.02 for business output per person and its shortfall from the factor model of 0.27 percentage points, combined with the negative shortfall in the employment rate of 0.42, yields an adjusted estimate of 1.09 (= 0.42×2.02 + 0.21) of the shortfall in business output per person from the factor model, adjusted for the fact that factor model underpredicts employment. This is larger than, but roughly comparable to, the sum of the irregular component computed using Okun's Law and the forecast error associated with the trend growth rate, which together add to 0.91. As another example, while the factor model overpredicts the average growth rate of the capital-output ratio (see Table 4), this ratio is countercyclical, and its growth rate exceeds the factor model forecast after adjustment the forecast for employment.

In Table 3 and Table 4, the contribution of the participation rate to the shortfall is the same. However, because the factor model underpredicts the recovery in the labor market, the contributions of other variables differ. For example, the contribution of TFP (not share weighted) is larger for the factor model than for Okun's law. These differences can largely be reconciled by the factor model's underprediction of the decline in the unemployment rate, combined with the Okun's law coefficient from Table 2. Put differently, the differences between Table 3 and Table 4 arise because the questions being asked in the two tables differ: in Table 3, what explains the slow growth in output, given the strength of the recovery in the unemployment rate? In Table 4, what explains the unusual aspects of both output and the labor market, given the state of the economy in 2009? The answer to the latter question involves unexpected improvements in hours, employment, and the unemployment rate, in addition to the variables that explain the divergence between the unemployment rate and output, namely TFP and participation.

In summary, this section documents that slow growth since 2009 is essentially entirely accounted for by slow TFP growth and declining participation. The crucial issue for interpreting these results is the extent to which the slowdown in TFP and the fall of participation were independent of, or alternatively a consequence of, the recession and its aftermath. For example, the financial crisis might have reduced innovative activity, thereby slowing TFP growth in a way not captured by our cyclical controls. More broadly, persistent headwinds to desired spending might have endogenously reduced the level or even growth in productivity; and hysteresis effects in labor markets might have reduced participation.

IV Why Have Capital Accumulation and Productivity Fallen Short?

We now turn to a closer examination of the related questions of the sources of the decline in productivity growth and the evidence on whether there has been an unusual hiatus in capital deepening—a reversal toward capital shallowing.

We have three findings. First, the decline in productivity growth has its roots before the recession. Evidence of this slowdown appears both in the aggregate quarterly time series data on productivity and annual data at the industry level.

Second, weak investment and capital growth does not appear to have been an important independent contributor to weak output growth. Growth of investment and of capital has been historically slow during the recovery, which on its face suggests that a source of the labor productivity slowdown is the lack of normal capital deepening. But this argument neglects the fall of the underlying growth rate of output from the decline in cyclically adjusted productivity and demographic (and perhaps other) changes that, as Section V concludes, have led to a secular decline in the labor-force participation rate. While capital formation has been slow, on net it has been no slower than output growth: by 2016, the capital/output ratio was in line with its long-term trend. Although deficiencies in business investment might have affected the timing of the recovery, after seven years capital was in line with historical norms.

Third, we find that most of the productivity slowdown occurred prior to the crisis. This is important because our first two conclusions point to the persistent fall in productivity growth as a key to understanding the slow recovery. If productivity slowed because of the weak recovery itself, for reasons not captured in our cyclical adjustment, then growth might pick up simply because the economy has returned to full employment. Our conclusion is that the slowdown reflects a pause, if not an end, to the broad-based, transformative effects of information technology. In particular, productivity growth was unusually high in the late 1990s and early 2000s in both the production of information-technology products and in the use of those products in other sectors, coinciding with the launch of the public internet and the proliferation of key technologies, such as the relational database.

IV.A When Did Productivity Growth Slow?

Even before the financial crisis, professional forecasters had noticeably lowered estimates of trend growth in labor productivity. Figure 10 plots the median forecasts from the Survey of Professional Forecasters for labor productivity growth over the next 10 years. The forecasts broadly track the lagging 10-year average growth of actual labor productivity computed using both real-time and finally revised data. Forecasts rose sharply between 1999 and 2000. They remained close to 2.5 percent through the 2006 survey. They have since fallen by about a percentage point. Half the decline in the forecasts occurred before the financial crisis, between 2006 and 2008.

The slowdown is also evident in the time-series data on TFP growth. Figure 6a shows that growth picked up in the mid-1990s and slowed prior to the recession. The statistical characterization of that change is an open question. It could be persistent change, or it could be transient good luck.

With respect to the timing and persistence of the slowdown in productivity growth, we undertake two sets of analyses. The first, frequentist in nature, entails computing tests for a break or for slower time variation in the mean of cyclically-adjusted productivity growth. The second, a Bayesian approach, provides posterior inference on whether the decline in the mean occurred before the 2007-2009 recession began.

Table 6 summarizes five tests for the null hypothesis that there is no time variation in the mean growth rate of TFP. Let y_t^{ca} denote the cyclically adjusted growth rate of productivity, so that, following equation (6), $y_t^{ca} = \mu_t + z_t$, where μ_t is the local mean (or trend) value of y_t^{ca}, and z_t is the mean-zero irregular component. The table shows results for two sample periods, a 60-year sample from 1956 through 2016 and the 35-year sample from 1981 through 2016,that has been the primary focus of this paper; here, we use the longer sample to increase power. The first three tests are the sup-Wald (the autocorrelation-robust Quandt Likelihood Ratio) break test of a constant mean against the alternative of, respectively, one, two, or three breaks. Along with the test statistic, this test yields estimates of the break dates themselves. The remaining two tests are the Nyblom (1989) tests that focuses power on small martingale variation in μ_t, and the LFST test (Müller and Watson (2008)), a low-frequency point-optimal test for martingale variation.

All five tests reject the null hypothesis that μ_t is constant using 1956 through 2016 sample, but not using the shorter 1981 through 2016 sample. In part this reflects increased power from the larger sample size, but also reflects the early 1970s productivity slowdown included in the first sample. Notably, the three-break full-sample test identifies break dates in 1973, 1995, and 2006, with a p-value (for the null of no breaks) of 0.01. These break dates accord with the conventional view of a high-growth period before 1973, a lower growth period until 1995, and the high growth period of the tech boom. Notably for our purposes, this boom is dated as ending before the 2007-2009 recession.

To gain additional insight into possible persistent changes in productivity growth, we adopt a latent variable state-space model for the trend and irregular components μ_t and z_t, in which μ_t is modeled as a Gaussian random walk and z_t is modeled as Gaussian white noise. By adopting a Bayesian framework, we are able to provide complementary insights into the timing of a peak in trend productivity growth and the magnitude of its decline prior to the recession. Details are given in the Econometric Appendix. This approach yields Bayesian posterior sets for μ_t that incorporate the uncertainty in the variance of $\Delta\mu_t$.

The results are summarized in Figure 11 and Figure 12. Figure 11 shows the 4-quarter growth rate of productivity, and three different estimates of μ_t: the cyclically adjusted biweight estimate, the three-regime estimate based on the estimated break dates in 1995 and 2006, and a 67 percent posterior interval for μ_t from the Bayesian implementation of the random walk-plus-noise model. Figure 12 provides the posterior distribution of the date of the maximum of the local mean of productivity growth between 1981 and 2016.

Taken together, we interpret Table 6, Figure 11, and Figure 12 as providing coherent evidence that the decline in productivity growth started before the recession. The posterior distribution in Figure 12 dates the peak of μ_t in the late 1990s or early 2000s, with little of the mass after 2006. The frequentist break tests in estimate a break date in 2006. Using the Bayesian approach, we can compute the posterior probability of the magnitude of the decline between the peak of μ_t around 2000 and its value in 2007: this calculation yields a posterior median estimate of 0.72 percentage points using the full sample, and a 67 percent posterior set of (0.32, 1.27).

These estimates, which suggest a significant decline prior to the cyclical peak, are also consistent with the decline in the biweight estimate and the Bayes posterior sets in Figure 11.

Fernald (2015) discusses other evidence of a pre-recession slowdown in productivity growth, especially from industry data. For example, the pre-recession slowdown was broad-based across industries—it was not, for example, particularly pronounced in housing-related industries or finance.

The discussion above focuses on measured productivity growth. A complementary perspective comes from looking at inputs to innovation, where a change in trend is apparent around 2000 or so—so even earlier than the shift for productivity.

In particular, productivity grows as the economy accumulates better ways to produce output. Some of the flows into the process of innovation and improvement are measured in the national income and product accounts. Figure 13 shows the log of the index of intellectual property investment from the accounts. It includes computer software, research and development spending in businesses, research at universities and nonprofits, and the production of books, movies, TV shows, and music. It is worth noting that the real growth rate of this category is 6.5 percent per year, far above the growth rate of any of the other series in this paper.

The graph shows that intellectual property investment grew faster than normal during the period of high productivity growth, grew more slowly than normal until the mid-1970s, and then entered a long period of high growth that came to an abrupt end in 2000 when the stock-market values of tech companies collapsed. Since 2000, IP investment has grown much more slowly than normal. The financial crisis in 2008 only slightly worsened the rate of contraction of IP investment relative to trend. The recovery that began in the economy as a whole in 2010 has so far done nothing to halt the low growth of investment in improved productivity. Recent research has attempted to measure additional intangible investments in innovation, training, reorganizations, and the like that are not currently included in the national accounts. Estimates of these additional intangible investments from Corrado and Jäger (2015) also show a slower pace of growth after about 2000.

The evidence on spending on innovation (as measured in the national accounts) thus also shows a slowdown much earlier than the recession. It is plausible that this spending might show up in measured productivity somewhat later, though the link need not be causal—both, for example, could be a reflection of the availability of ideas or other factors. In U.S. data since the early 1970s, the unusual period for productivity growth was the decade from 1995 to 2005.

IV.B Why Has Capital Fallen Short?

On its face, concerns about weak investment seem appropriate. Figure 14 shows the log of real business investment in equipment since 1984. This form of investment is a major fraction of capital formation and embodies many of the new technologies that account for productivity growth. The most prominent feature of this series is its rapid growth in the 1990s. The tech collapse in 2000 resulted in a relatively small contraction followed by expansion in the mid-2000s. Equipment investment was well above trend in 2007 and even a bit above trend in 2008. It fell almost in half (just under 0.5 log points) in 2009, a much larger percentage drop than in any previous recession in the years since 1948. Equipment growth since 2000 has been lower than in the 15 earlier years.

The shortfall in capital formation could reflect many factors, some of which may be tied to special features of the recession and recovery, though others are more general. These include tight credit for some borrowers, increased financial frictions, heightened uncertainty, regulatory barriers, increased market power, or other factors. To assess these ideas, we need a model of capital formation. The core of such a model is a demand function for productive capacity. That demand is derived from the demand for output, and also depends on the cost of financing—as laid out by Jorgenson, with Tobin's addition of adjustment costs.

An implication of investment theory is that if investment were an important, independent factor explaining the weak recovery, then capital's contribution to the labor-productivity growth accounting in equation (5) should show large deviations from previous experience. But, as discussed already, calculations based on that equation assigns essentially no role to a capital shortfall: the capital-output ratio has been completely in line with its modest upward trend. In other words, capital growth has been weak, and output growth has been weak, but after cyclical adjustment, the ratio of the two has behaved normally—see Figure 6(d).

Another way to view investment is through the return to capital. An important determinant of business investment is the payoff to owners of capital. Some accounts of weak investment imply that capital was not earning as much as in normal times. But, as Figure 15 shows, the earnings of capital, measured as the sum of business profits, interest paid, and depreciation, have been remarkably steady since the crisis. Earnings per dollar of capital fell in 2009, but rebounded to normal in 2010 and have remained normal since.

More broadly, investment dynamics in recent decades are complex and the pattern across industries is nuanced. As Alexander and Eberly (2016) and Gutierrez and Phiippon (2017) highlight, the apparent weakness in investment started around 2000, not with the 2007-09 recession. So this is yet another example of a trend whose origins predate the crisis.

Gutierrez and Philippon, in particular, find that investment has been weak despite high Q, which they attribute to weak competition and governance changes. Both of these stories would tend to predict that the steady-state capital-output ratio should fall. In contrast, as noted earlier, a decline in expected growth should tend to raise that steady-state ratio. It is beyond the scope of our paper to sort out quantitatively which effect dominates for the capital-output ratio. Our empirical evidence that the capital-output ratio is on its previous trend is consistent with the two forces roughly offsetting. In any case, in terms of investment, they point in the same direction--investment should, at least for a time, be unusually weak for reasons unrelated to the recession or slow recovery per se.

IV.C Explanations for Slow Productivity Growth

Why has productivity growth been so slow if it's not the result of the financial crisis? Our conclusion is that the slowdown is a pause in—if not an end to—the information-technology revolution. Our related conclusion is that the slowdown was not mainly the result of the recession. In this section, we review a variety of hypotheses about the productivity slowdown. We begin with three non-recession explanations.

1. *Mismeasurement*. Perhaps the problem of slow growth in both productivity and output is illusory? That is, perhaps we aren't fully measuring the gains from tech-related hardware, software, and digital services? This hypothesis, if true, would undercut the entire motivation for this paper. Subjectively, IT-related innovation still feels rapid to many people—after all, we can all do amazing things on our phones today that we couldn't do in 2005. Of course, mismeasurement concerns aren't new. And for mismeasurement to explain the productivity slowdown, growth must be mismeasured by more than in the past.

In this regard, Byrne, Fernald, and Reinsdorf (2016) and Syverson (2016) find no evidence that, on balance, the mismeasurement of the growth rate of tech-related real output has gotten *worse* since the early 2000s. We have always had mismeasurement. Moreover, the steady shift of economic activity towards poorly measured services, such as health care, also does not change the picture. Measured productivity growth in these sectors always been low, but the mid-2000s slowdown in productivity growth spread broadly across industries. Thus, changes in weighting matter relatively little. Aghion and others (2017) find a modest increase after the early 2000s in missing growth from creative destruction and increases in varieties. But the increase in bias is small relative to the measured slowdown in productivity growth.

2. *Rising regulation and loss of dynamism*. Some commentators have pointed to a rising regulatory burden as a potential reason for slowing productivity growth (Barro, 2016). In cross-country contexts, differing regulatory barriers do seem to matter (Fatas 2016). Cette, Fernald, and Mojon (2016) compare the gains from information technology in the U.S. and Europe.

Europe didn't get the same productivity benefits from tech after 1995 as the U.S. did. The leading hypothesis is labor and product-market inflexibilities—many induced by regulations—that limited the ability of firms in Europe to reorganize to benefit from tech investments.

Rising regulation could be a reason for the observation that, by many measures, dynamism in the U.S. economy has declined since the 1980s (Decker and others, 2016a, b). Job creation and destruction has slowed; the business startup rate has fallen; and young firms have grown less in recent years.

In the U.S., a rising Federal regulatory burden does not appear to explain the medium-frequency variations in productivity. First, although some commentators have pointed specifically to post-2008 regulatory changes, the timing does not fit because the peak in productivity growth occurred well before that time. With the exception of the decade starting in 1995, relatively slow productivity growth has been the norm since the 1970s.

Second, even for the post-2008 period, the industries where regulation increased the most did not for the most part show a decline in productivity growth. Al-Ubaydli and McLaughlin (2015) applied text-analysis methods to the *Code of Federal Regulations* to construct industry-level indices of regulations from 1970 through 2014. Their "RegData" database covers 42 private industries matched to Bureau of Labor Statistics' industry-level productivity data for the private business economy, which runs from 1987 through 2014. These data are described further in the appendix. The industries that saw greatest increases in regulation after 2008, compared with growth rates from 2000 to 2008 were, most notably, (i) finance (credit intermediation, funds and trusts, securities, and insurance), (ii) energy (pipelines, oil and gas extraction, and utilities, especially); (iii) construction, and (iv) transportation (especially trucking, water, and rail).

Table 7 presents selected cuts of the industry productivity data on the growth rate of business productivity. The slowdown for the entire private business economy (line 1) after 2004 is even more pronounced in these data than in the Fernald data. Finance slows sharply after 2004 and shows no further slowdown after 2007, the period of Dodd-Frank and other restrictions. The energy industry experienced *faster* productivity growth after 2007, reflecting the fracking revolution. Energy regulations are certainly not the reason for the broadbased productivity slowdown. Construction also has experienced less negative productivity growth. Of heavily regulated industries, only transportation has had lower productivity growth, but it is only 2.5 percent of value added.

Perhaps finance matters because of its importance as an intermediate provider of services. Using the input-output tables, we divided industries into finance-intensive (row 8) and non-finance-intensive (row 9) industries, defined as expenditure on financial services relative to industry gross output. Both groups slow sharply after 2004, but the finance-intensive grouping actually improved after 2007, when finance restrictions tightened. Over the entire post-2004 period, the slowdown is larger for non-financial-intensive industries. Thus, it does not appear that post-2008 financial restrictions were a major impediment to productivity growth.

Third, there is little evidence of a broader regulatory effect. Table 8 shows panel regressions of industry productivity growth on current and lagged values of growth in industry regulatory restrictions. All regressions include industry fixed effects; the second column includes year effects. Columns (1) and (2) show that, with one and two lags, growth in regulatory restrictions does retard productivity growth. But the effect is never statistically significant, and the explanatory power is tiny. Columns (3) and (4) try averaging lagged values, but these also yield small and statistically insignificant effects.

This finding is consistent with Goldschlag and Tabarrok (2014), who find that changes in U.S. federal regulations have little or no effect on industry entrepreneurial activity or dynamism. The lags may be long and uncertain—and thus hard to detect—or that the regulations that matter are mainly at the state and local level. A commonly held view is that some or many of these regulations—such as overly restrictive land-use restrictions and onerous occupational licensing—constrain activity. But, at the macroeconomic level, we cannot find evidence that regulation is a first-order issue for explaining recent slow productivity growth.

More broadly, there is a question of whether declining dynamism is an independent contributor to the slowdown. Decker and others (2016b) suggest that the character of declining dynamism changed after 2000, which would match the view that there were structural shifts in trend growth prior to the 2007-2009 recession. The direction of causation between innovation and dynamism is also not necessarily clear-cut. For example, if the shortfall is the lack of available or exploitable ideas for the broad economy, then the lack of dynamism might be a symptom of that lack of opportunity.

3. *A pause in the information technology revolution*. The hypothesis that tech was the culprit is natural. A large literature links the mid-1990s speedup in productivity growth to the exceptional contribution of computers, communications equipment, software, and the Internet. The idea is that tech has had a broad-based and pervasive effect on the economy through its role as a general purpose technology (Bresnahan and Trajtenberg, 1995; David and Wright, 2003; Basu, Fernald, Oulton, and Srinivasan, 2004). That is, it fosters complementary innovations, such as business reorganization to take advantage of an improved ability to manage information and communications. Businesses throughout the economy transformed how they operated and became more efficient. But, by the early 2000s, industries like retailing had already been substantially reorganized, after which the gains from further innovation might have been more incremental than transformative (Gordon, 2016; Fernald, 2015).

Table 7 suggests some evidence consistent with this hypothesis. Tech-producing industries (line 5) grew much slower after 2000 and even slower after 2007. Industries that use tech intensively show a larger slowdown after 2007 relative to the period from 2000 through 2004. But it is fair to say that the slowdown is broad-based. All industries use tech, and increasingly so. If that is the story, we might see another such period in the future, perhaps reflecting artificial intelligence, cloud computing, the Internet of things, and the radical increase in mobility from smartphones. We have not yet seen those gains in the data.

This story rings true in a number of ways. First, it is consistent with the large literature on the role of tech in the productivity acceleration in the late 1990s. Second, it is consistent with the view in the general-purpose-technology literature that the gains are, essentially, a series of drawn-out levels effects. It is hard to predict how long the gains will continue. The gains might ebb and flow for a time (Syverson, 2013).

4. *Fallout from the recession and financial crisis*. Our use of cyclically adjusted productivity growth corrects for normal cyclical movements in productivity and in particular allows us to focus on the magnitude and timing of the more persistent, secular slowdown that has been the focus of this section so far. But were there special features of the 2007 recession, such as its origins in the financial crisis and its depth, that contributed to the slowdown?

Theory is ambiguous about whether severe recessions, including financial ones, have a persistent effect on the path of productivity—both its growth rate and its level. A crisis could reduce the invention or adoption of new technologies (Fatas, 2000, 2002; Reifschneider, Wascher, and Wilcox, 2013; Anzoategui, Comin, Gertler, and Martinez, 2016). Liu and Wang (2013) model a financial accelerator that leads to procyclical reallocation and productivity. Sedlacek and Sterk (2013) find that not only did the number of U.S. startups drop sharply during the 2007-2009 recession, but that firms born during recessions tend to be smaller and less productive than others even after the economy recovers. If weak productivity growth were primarily a result of the recession and slow recovery itself, then a high-pressure economy might help reverse those effects and lead to faster growth in innovation and technology (Yellen, 2016).

Theory reaches ambiguous conclusions. Reallocation effects in some models go the other way, raising measured productivity in a credit crisis (Petrosky-Nadeau, 2013), or the cleansing effects described by Caballero and Hammour, 1994). Bloom (2013) points out that higher uncertainty can stimulate longer-run innovation.

Overall, there is limited empirical evidence for developed countries that historical business-cycle downturns, financially related or otherwise, permanently cut the level or growth rate of productivity. The depressed 1930s were, by all accounts, an extraordinarily innovative period (Field, 2003, Alexopoulos and Cohen, 2011, and Gordon, 2016). Oulton and Sebastiá-Barriel (2014) perform growth-accounting exercises across countries following financial crises. For developed countries—but not for others—the long-run level of productivity is essentially unchanged by a financial crisis—indeed, the point estimate is slightly positive. For the U.S., Huang, Luo, and Starts (2016) find that the level of productivity bounced back quickly from recessions, including after 2009. It is unclear it is a major factor for the U.S. relative to the pre-recession slowdown. Nor is it the entire story for continental Europe, where productivity has diverged from U.S. levels since the mid-1990s (Cette and others, 2016).

The biggest challenge for explaining U.S. data is the timing. Productivity growth slowed prior to the recession. Anzoategui and others (2016) argue that there was a pre-recession shock to exogenous growth combined with the large shock from the recession. As noted in Section IV, there is limited U.S. evidence that investments in research and development and other intellectual property slowed because of the recession. Rather, the pace of growth slowed earlier.

We conclude that it is difficult to measure counterfactual productivity growth absent the recession, or absent the regulatory tightening. But we find that the weight of the evidence suggests that the slow pace since the mid-2000s is real, contributed substantially to the disappointing recovery, and may well continue.

V Changes in the Labor market

Table 3 shows that the contribution of participation in the labor force to output, after a tiny cyclical adjustment declined at a rate of 0.69 percentage points per year (column e), compared to an increase of 0.15 points per year averaged over the three previous recoveries (column d), for a shortfall of 0.85 points per year (column f). Cumulated over the recovery through 2016, the shortfall was 6.11 percentage points (column i), almost as large as for TFP.

Prior to the crisis, recessions only slightly depressed participation—unemployment rose by almost the same amount that employment fell. With higher unemployment, participation was discouraged by the added time needed to find a job. But wealth and income fall in recessions. The loss induces more people to seek and take jobs, and so is a force that raises participation. In previous recessions, the two forces approximately offset each other. The cyclical coefficient in Table 2 is -0.16 over a sample period that includes the rise in unemployment and fall in participation during and after the recession. Although this generalized Okun's coefficient increases with the addition of more lags of the unemployment rate, which allows for longer dynamic adjustment of the participation rate, even with two years of lags it is only -0.27. The estimate declines in magnitude with additional lags. Regardless of the lag specification, by 2016 the normal cyclical component of the participation rate was essentially zero. As a result, our analysis points to the fall in participation from the trough to 2016 as entirely due to a fall in the trend and irregular parts.

The labor force comprises people 16 and over who are working or are actively looking for work. Over the past 50 years, trends in participation have been quite different for men and women, so we consider them separately. Figure 16 shows the percentages of men and women in the labor force starting in 2006. Though the rates moved divergently in earlier years, with women rising and men falling, the two rates moved together from 2006 onward. Both declined substantially after 2008. Determining the counterfactual—what would have happened to participation had the trauma of 2008 and the long slump following not occurred?—is a challenge. We find that, although some forces determined long before the crisis depressed participation, other forces specific to the post-crisis years account for about two-thirds of the decline during the recovery from 2010 through 2016.

Many authors have ascribed part of the decline in participation to demography, specifically to the rising fraction of the population aged 55 and above. Traditionally, this age group tends to exit the labor force through retirement. But adjusting for age composition alone misses some demographic forces that reduce the propensity to retire. In particular, the people who moved into the 55-plus age group during the recovery are better educated than their predecessors, as they belong to cohorts that were more likely to finish high school and attend college. Calculations of pure-aging declines in participation, which use historical rates for older workers, could overstate the contribution of aging during the recent recovery because those better-educated, now-older workers would normally retire later in life. Accordingly, we calculate indexes that adjust for five demographic dimensions of heterogeneity in the working-age population.

The measured overall labor-force participation rate can be written as

$$L = \sum_i s_i L_i \ , \tag{11}$$

where s_i is the population share and L_i is the participation rate of demographic group i. The change in the overall participation rate satisfies

$$\Delta L = \sum_i s_i \Delta L_i + \sum_i L_i \Delta s_i \tag{12}$$

to a high degree of accuracy, especially if s_i in the first term and L_i in the second are measured as equally weighted values from the earlier and current periods. The cumulation of the first term is the component of the level of participation attributable to changes in participation within demographic groups and the cumulation of the second term is the component attributable to composition changes in the population. We call these the *rate* and *share* effects. Indexes calculated this way are named after Divisia and the refinement of measuring shares as equally weighted averages is named after Tørnqvist. The variation in the rates over the period is high enough to make any share index with fixed rates misleading. Counterfactual calculations based on holding rates at, say, the 2006 or 2016 levels are effectively fixed-rate indexes.

We have implemented this approach with annual data from the CPS for about 6,100 detailed cells defined by 67 age categories; two sexes; four education groups; four race groups; and three marital status groups. A few hundred of the cells in each year are empty. Figure 17 shows the overall participation rate and our *rate* index. Because the residual in the index calculation is tiny, the difference between the two indexes is effectively our index of the *share* effect, that is, the effect of changing demographics. During the recovery, from 2010 through 2016, the reported participation rate, across the population aged 16 and older, fell by 1.8 percentage points. Of this, 1.2 points came from the rate effect—the result of lower participation, on average, within demographic groups—and 0.6 points came from compositional change. In other words, forces other than demography accounted for about two-thirds of the overall decline during the recovery, and for about one-half of the decline since the cyclical peak in the fourth quarter of 2007.

The key question is, what are the reasons for the large non-demographic decline in the participation rate?

One possibility is that our cyclical adjustment methods are flawed, and that there is a large cyclical component of the participation rate that will ultimately fade away as long as the labor market remains reasonably tight. This argument would be consistent with some authors who have argued for a large cyclical component in the decline in the participation rate. Erceg and Levin (2014) use state-level data to study the relation between unemployment and participation. Their model has the change in participation, in percentage points, between 2007 and 2012, as the left-hand variable and the change in unemployment between 2007 and 2010 as the right-hand variable. The estimated coefficient is -0.30 in their preferred specification (their Table 2, p. 12). As discussed above, our estimate of the cyclical component is sensitive to the number of lags in the time series regression; with 12 lags, to match Erceg and Levin's specification, the generalized Okun's law coefficient is -0.19. But by the middle of 2016, the unemployment rate had long since peaked and had returned to a normal or near-normal range, and it stabilized around 4.7 percent. A large cyclical coefficient suggests that cyclical factors played a role in the decline in participation early in the recovery, but by 2016, even a large cyclical coefficient implies only a very small normal cyclical component by mid-2016.

If the non-demographic participation gap as of 2016 is not part of a normal cyclical pattern, it must either be a response to an unusual feature of this recession and recovery, or the continuation of a phenomenon that began before the recession. While this recession was certainly large, the 5.5 percentage point increase in the unemployment rate from its 2006 trough to its 2009 peak was comparable to the 5 percentage point increase from its 1979 trough to its 1982 peak spanning the twin recessions of the early 1980s. Peak unemployment in 2010 was less than its peak in 1982. As shown in Figure 1, the recovery of the unemployment rate in the current recession was comparably fast to its recovery in the early 1980s. Because the cyclical movements of the early 1980s are part of the data set used to estimate the Okun's law coefficients, explanations that appeal to hysteresis must therefore argue that the correlations from previous cycles do not translate to the current cycle. It is not possible to estimate these coefficients precisely using only the current cycle; but, if anything, the unemployment coefficients are smaller when the current cycle is included in the data set. Finally, a related critique is that the coefficients in the generalized Okun relation are different for increasing than decreasing rates of unemployment, so that our cyclical estimate is mis-specified; but we examined this empirically too and found no evidence of this interaction effect. To put the point in a different way, the hypothesis that rising unemployment discouraged participation had some support in the contraction, but failed to describe the relation between unemployment and participation in the recovery.

Aaronson and co-authors (2014) report a wide variety of results on participation. They find that their forecasts of participation published in 2006 were remarkably accurate as of 2004, suggesting that the entirely unforeseen recession and recovery that began at the end of 2007 had little net effect on participation. Their overall conclusion is that the sources of the decline in participation are partly demographic and partly a change not much related to conditions in the labor market. Though they do not discuss the expansion period beginning in mid-2009 specifically, it appears that their results confirm our conclusion that the dramatic improvement in the labor market during the recovery had little net effect on participation. They cite a number of studies of participation with similar conclusions.

Our conclusion is that the roots of the non-demographic participation gap as of 2016 lie somewhere other than in the recession. While there has been increasing research interest in what these roots are, that work has so far been inconclusive.

Figure 18 provides additional information useful in trying to understand the decline in participation. It shows participation rates for people aged 25 through 54, broken down by family income. Between 2004 and 2013, participation *rose* among members of the poorer half of families, and *fell* substantially in the upper half, the third and fourth quartiles. Essentially all the decline in participation occurred in families with higher incomes. This finding points away from the hypothesis that the decline in participation represented marginalization of poorer families from the labor market.

Table 9 investigates how people spent the time freed up by reduced work and job search. It compares time allocations in 2015 to 2007. Market work, including job search, fell by 1.6 hours per week for men and by 1.4 hours for women. The two categories with increases were personal care and leisure, which includes a large amount of TV and other video-based entertainment, especially for men. The decline in hours devoted to other activities included a decline in housework for women. Basically, time use shifted toward enjoyment and away for work-type and investment activities. There was no substitution from market work to either non-market work or investment in human and household capital.

The surprising, large, and persistent decline in labor-force participation is a phenomenon that deserves and will receive intensive study. While there is room for disagreement about the extent to which the decline in participation during the early recovery was a response to an extremely slack labor market, that cyclical component was gone by mid-2016. Similarly, although demographic shifts are and will continue to be an important part of the decline in the participation rate, the idea that this decline is mainly the result of demographic shifts has also not held up. The successful explanation will consider changes in family structure, real wages, taxes, benefits, and the value of time spent outside the labor market, along with the tightness of the labor market.

VI Other Explanations for Slow Output Growth

So far, our discussion has focused on understanding the recovery using the growth accounting decomposition. While we believe this decomposition is central to understanding the recovery dynamics, including those that stress long-term demographic changes, it does not directly address a large number of proposed explanations for weak growth. We therefore now turn to some of those other explanations. Our forecasting model provides evidence about some of these ideas.

Before considering the ideas individually, we note that our earlier results take demand into consideration through the use of unemployment as a cyclical indicator, and through the use of a factor model with a multivariate statistical characterization of the cycle. If we are correct that unemployment is a good statistical indicator and that unemployment rates below five percent imply an economy in a cyclically normal condition, then explanations based on the persistence of weak demand are ruled out. Moreover, explanations based on demand deficiency need to reconcile them to the fact that the recovery of the unemployment rate that was as fast or faster than normal. Sponsors of explanations based on weak demand need to couple their explanations with a parallel explanation of the behavior of labor-market indicators during the recovery.

We also consider explanations suggests that at least some of the slowdown in productivity and output growth is an artifact of escalating challenges in the measurement of real output and prices in some important sectors of the economy, such as information technology.

VI.A Empirical Evidence from the Forecasting Exercise

Figure 7, Figure 8, and Figure 9 show three periods in the history of the recovery. From mid-2009 through 2010, the economy grew vigorously, with employment, output, consumption, and private fixed investment, all growing at or above the forecast path. From 2011 through 2013, although employment growth was strong, it was below its predicted path, and the associated predicted strong growth in output failed to materialize. This period had a large growth gap—it lacked the sustained output growth in the 3 to 4 percent range typical of earlier recoveries. After an initial surge in 2009, the growth of productivity was low during this period, well below its predicted path—see Figure 7(i). In the third period, since 2014, growth in many aggregates, including output and especially employment, has been stronger than the forecast path, and—notably—the slow productivity growth over this period is consistent with the cyclical prediction. The picture is one of a recovery delayed: the slow-growth puzzle is largely the absence of strong growth in productivity and output in 2011 through 2013.

The demand decomposition in Table 5 indicates that most of the demand components tracked their forecast paths on average. Although exports were unexpectedly weak, so were imports, after share-weighting their contributions to the average shortfall in output growth is negligible, 0.03 and -0.01 percentage points per year, respectively. Table 5 indicates that the average forecast error is largely attributable to three sources: consumption of services (0.18 percentage points), federal government expenditures (0.20), and state and local government expenditures (0.12).

For federal government purchases, the main shortfall occurred in 2013 and 2014. This period coincides with the fiscal drag associated with unwinding Recovery Act expenditures and with the sequester. For state and local expenditures, the period of negative contributions was longer, from 2010 through early 2014.

Consumption growth over the recovery was slightly weaker than predicted—a 0.26 percentage-point contribution to the output shortfall. Most of this weakness is attributable to two service sectors: housing and utilities (0.07 percentage points) and financial services and insurance (0.07 percentage points).

The forecast error in residential investment averaged -0.09 percentage points over the full period, but this masks the delayed recovery in the housing sector. Through 2011, the normal cyclical recovery in housing did not materialize, and housing investment growth did not stabilize around the forecast path until 2012. The strength of the housing market since 2014 accounts for the negative contribution of residential investment to the output shortfall. .

VI.B Discussion

Because we do not identify a structural factor model, we do not identify the structural shocks that led to the 0.57 percentage points of slow GDP growth over this period. Nevertheless, the pattern of forecast errors sheds light on some of the explanations for the slow recovery.

Explanations in which aggregate demand is held back by unusually retarded growth of consumption—increasing inequality, policy uncertainty, or consumer deleveraging—do not square with the fact that contribution of consumption growth to the shortfall in output growth was only 0.26 percentage points; rather, consumption growth largely tracked its predicted path over the recovery. Moreover, the largest shortfall in consumption is in services, mainly housing services and financial services and insurance, and in the latter case, for only three aberrant quarters in 2011 and 2012. This pattern does not seem to align with any explanation that focuses on shortfalls in aggregate demand that operate through consumption broadly.

Similarly, there is little evidence to support theories that operate through unduly slow investment. Nonresidential investment growth was, in fact, unexpectedly strong early in the recovery, and otherwise largely tracked its predicted path, except for a slow spell in 2013 (Figure 8(h)).

The fact that the growth of consumption and investment largely tracked their historical cyclical patterns suggests that unusual features of the current recession that held back the normal cyclical growth of aggregate demand are not key drivers of the slow recovery. Moreover, one would expect slow aggregate demand to be reflected in sluggish revival of employment and the unemployment rate, but that is evidently not the case because employment growth exceeded the 2009 prediction on average. Growth was strong early and late in the recovery. One nonstandard explanation that has circumstantial support is that there has been hysteresis in the labor market, with an unusually prolonged recovery of the long-term unemployment rate and the shortfall of participation exceeding the combined predicted effects of demographics and normal cyclical patterns.

Our examination of the components of demand does show one unusual feature of the evolution of demand that made a contribution to the slow recovery: the weakness in both federal and state and local government purchases. The timing of the forecast errors suggests that the unwinding of the Recovery Act spending combined with the sequester provided substantial headwinds to the recovery, an estimated 0.20 percentage points of reduction in mean growth over this period, relative to the predicted path. In addition, the persistently slow growth of state and local government purchases through 2013, along with the slow growth over this period of state and local government employment, points to unusually severe fiscal drag imparted by restrained state and local purchases associated with balanced budget requirements and the prolonged effect on real estate tax receipts of the fall in house prices during the recession. These measures do not include transfers. However, the addendum line to Table 5 adds government consumption and transfers. This category was growing, unlike direct government purchases, which were shrinking. So transfers may have somewhat supported consumption. Nevertheless, there was still a large shortfall. The DFM forecasts that this composite category should have grown 2.86 percent per year, but in fact it only grew at a 0.66 percent pace. .

Finally, we find some room for explanations associated with poor or missed measurement of real output. Gross domestic income growth averaged 2.34 percent over 2009 through 2016, while GDP grew at 2.06 percent. Table 5 suggests that some of this difference may come from unexpected sources. In particular, half of the unexpected decline in services consumption in 2013 is attributable (in a national accounting sense) to a decline in one of the most poorly measured sectors of consumption: financial services and insurance. Additional investigation of these measurement issues is warranted.

VII Concluding Remarks

Output grew substantially less in the recovery from the 2007-2009 recession than would normally have accompanied the healthy decline in unemployment. It grew less than it would have given its normal relation to an index derived from many macro indicators. And it grew less than had been forecasted at the time of the trough in mid-2009. An explanation for poor output growth needs to start with two key facts—productivity grew substantially less than its historical growth rate, both in expansions and in general. And labor-force participation shrank an atypical and unexpected amount. Research on both topics is active today. We conclude in this paper that the large movements in both factors were in train prior to the recession, and cyclical effects contributed at most modestly to them.

An important question is whether growth will pick up in the future, or slow further. For example, the median respondent in the *Survey of Professional Forecasters* for 2017, first quarter, forecasts growth in the next three years, and the next 10, to exceed its average pace over the recovery so far.

Although changes in technology trends are hard to predict, the analysis in our paper does not support such optimism. The disappointing average pace since 2009 included a large cyclical component that will go away. The remaining slow underlying pace of growth instead reflected underlying non-cyclical trends that predated the recession and that have been persistent, to date. Thus, the growth seen during the recovery might, for a while, be as good as it gets.

Data Appendix

Growth and expenditure-side decompositions of output

Our main growth-accounting data for the U.S. business sector are described in detail in Fernald (2014). Those data are available quarterly, in growth rates, from 1947:Q2 on at http://www.frbsf.org/economic-research/economists/jfernald/quarterly_productivity.xls. The version used in this paper were prepared on December 30, 2016.

For the overall economy, output is measured by real gross domestic product (GDP) and the geometric average of GDP and real gross domestic income (GDI) (see Nalewaik (2010), Greenaway-McGrevy (2011), and Aruoba et al (2012)). We refer to the average as gross domestic output (GDO). Business sector output is also GDO using Fernald's measure.

Per-person values are formed using the civilian noninstitutional population 16 years of age and older from the Bureau of Labor Statistics Current Population Survey (FRED series CNP16OV). Other BLS-CPS variables include employment (CE16OV), labor force (CLF16OV) and the civilian unemployment rate (UNRATE). Quarterly data were constructed by averaging the monthly data for each quarter.

The expenditure variables (Table 5) are from the Bureau of Economic Analysis N.I.P.A. accounts.

Industry level TFP, finance intensity, and regulation data,

Bureau of Labor Statistics multifactor productivity (MFP) data and industry capital data were downloaded from http://www.bls.gov/mfp/mprdload.htm (accessed September 6, 2016). Growth-rate data run 1988-2014. The industry classification system is NAICS. See the online appendix to Fernald (2015) for details on how the data were manipulated and aggregated.

IT intensity is based on factor shares, i.e., payments for IT as a share of income. "IT intensive" is the set of industries with the highest IT shares that constitute 50 percent of the value-added weight (averaged 1987-2014) for the business sector excluding finance and direct IT production. For finance intensity, we aggregated industries from annual BLS I-O tables (accessed February 23, 2017) from http://www.bls.gov/emp/ep_data_input_output_matrix.htm. The finance share was nominal purchases of intermediate financial services as a share of industry gross output. "Finance intensive" is set of business (excluding finance) industries with the highest finance shares constituting roughly half the value-added weight.

Al-Ubaydli and McLaughlin (2015) produced the regulation data, available at regdata.org. The website summarizes the data as "RegData is a database that quantifies the number of individual restrictions in the Code of Federal Regulations and…determines which industries are targeted by those regulatory restrictions." They match regulations to BEA industries, which we then matched with BLS industries. Not all industries have reliable measures of regulation, and those industries are omitted. The included industries cover more than 80 percent of private value added and, when aggregated, have a similar TFP pattern to overall private business.

Labor force participation rates

The data underlying the demographic and family income decompositions for labor force participation are from the CPS. [to be augmented] .

References

Aaronson, Stephanie, Bruce Fallick, Andrew Figura, Jonathan Pingle, and William Wascher. 2006. "The Recent Decline in the Labor Force Participation Rate and its Implications for Potential Labor Supply." *Brookings Papers on Economic Activity,* issue no.: 69-154.

Aaronson, Stephanie, Tomaz Cajner, Bruce Fallick, Felix Galbis-Reig, Christopher Smith, and William Wascher. 2014. "Labor Force Participation: Recent Developments and Future Prospects." *Brookings Papers on Economic Activity*, 45 (2 fall): 197-275.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li. 2017. "Missing Growth from Creative Destruction." Federal Reserve Bank of San Francisco Working Paper 2017-04. http://www.frbsf.org/economic-research/publications/working-papers/wp2017-04.pdf

Alexander, Lewis and Janice Eberly. 2016. "Investment Hollowing Out." Manuscript.

Al-Ubaydli, O. and McLaughlin, P.A. 2015. "RegData: A Numerical Database on Industry-Specific Regulations for all United States Industries and Federal Regulations, 1997-2012." *Regulation & Governance*, 11: 109–123.

Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez. 2016. "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence." NBER Working Paper No. 22005.

Barro, Robert J. 2016. "The Reasons Behind the Obama Non-Recovery." *The Wall Street Journal*, September 20. http://scholar.harvard.edu/files/barro/files/wsj_published_version_092116.pd_f

Basu, Susanto and John G. Fernald, 2001. "Why Is Productivity Procyclical? Why Do We Care?" in: *New Developments in Productivity Analysis*, 225-302. National Bureau of Economic Research, Inc.

Basu, Susanto, John G. Fernald, Nicholas Oulton, and Sally Srinivasan. 2003. "The Case of the Missing Productivity Growth: Or, Does Information Technology Explain Why Productivity Accelerated in the United States but Not the United Kingdom?" *NBER Macroeconomics Annual*.

Blanchard, Olivier, Guido Lorenzoni, and Jean-Paul L'Huillier. 2017. "Short-run Effects of Lower Productivity Growth: a Twist on the Secular Stagnation Hypothesis." NBER Working Paper 23160.

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

Bosler, Canyon, Mary C. Daly, John G. Fernald, Bart Hobijn. 2016. "The Outlook for U.S. Labor-Quality Growth." Federal Reserve Bank of San Francisco Working Paper 2016-14. http://www.frbsf.org/economic-research/publications/working-papers/wp2016-14.pdf.

Bresnahan, T. F., and M. Trajtenberg. 1995. "General-Purpose Technologies: Engines of Growth"? *Journal of Econometrics* 65(Special Issue, January): 83-108.

Byrne, David M., John G. Fernald, and Marshall B. Reinsdorft. 2016. "Does the United States have a Productivity Slowdown or a Measurement Problem?" *Brookings Papers on Economic Activity*, Fall: 109-157.

Caballero, Ricardo J. and Mohamad L. Hammour 1994. "The Cleansing Effect of Recessions." *American Economic Review*84(5): 1350-68.

Cette, Gilbert, John G. Fernald, and Benoît Mojon. 2016. "The Pre-Great Recession Slowdown in Productivity." *European Economic Review* 88(C): 3-20.

Congressional Budget Office (CBO) 2006. "The Budget and Economic Outlook: Fiscal Years 2007 to 2016." https://www.cbo.gov/publication/41661.

Congressional Budget Office. 2012. "What Accounts for the Slow Growth of the Economy After the Recession?" November 2014.

Congressional Budget Office. 2014. "The Slow Recovery of the Labor Market" *Congressional Budget Office*. February 2014.

Council of Economic Advisers. 2016. "The Long-Term Decline in Prime-Age Male Labor Force Participation." https://obamawhitehouse.archives.gov/sites/default/files/page/files/20160620_cea_primeage_male_lfp.pdf

Corrado, Carol, and Kirsten Jäger. 2015. "Wealth and Investment in Mature Societies." Presentation given at Milestone 3: Midterm Conference, Smart Public INTANgibles (SPINTAN), London, April 23–24. http://www.spintan.net/wp-content/uploads/public/Carol-Corrado-Wealth-and-Investment_carolcorrado.pdf

Council of Economic Advisers. 2015. "A Better Measure of Economic Growth: Gross Domestic Output (GDO)." *Council of Economic Advisers Issues Brief*, the White House, July 2015. Accessed January 28, 2017.https://obamawhitehouse.archives.gov/sites/default/files/docs/gdo_issue_brief_final.pdf.

David, Paul and Gavin Wright 2003. "General Purpose Technologies and Productivity Surges: Historical Reflections on the Future of the ICT Revolution." In Paul A. David and Mark Thomas (eds.), *The Economic Future in Historical Perspective*.

Daly, Mary C., John G. Fernald, Fernanda Nechio, and Oscar Jorda. 2017. "Shocks and Adjustments." FRBSF Manuscript.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2016. "Declining Business Dynamism: What We Know and the Way Forward." *American Economic* Review 106, no.5: 203–7.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2016. "Where Has All the Skewness Gone? The Decline in High-Growth (Young) Firms in the U.S." *European Economic Review* 86: 4-23

Elsby, Michael W.L., Bart Hobijn, and Ay¸segul Sahin. 2015. "On the Importance of the Participation Margin for Labor Market Fluctuations," *Journal of Monetary Economics*72: 64--82.

Erceg, Christopher J. and Andrew T. Levin. 2014. “Labor Force Participation and Monetary Policy in the Wake of the Great Recession,” *Journal of Money, Credit and Banking* 46 (S2): 3–49.

Fatas, Antonio. 2000. “Do business Cycles Cast Long Shadows? Short-Run Persistence and Economic Growth.” *Journal of Economic Growth* 5:147-162.

Fatas, Antonio. 2002. "The Effects of Business Cycles on Growth." *Economic Growth: Sources, Trends and Cycles*. Eds. Norman Loayza and Raimundo Soto. Central Bank of Chile.

Fatas, Antonio. 2016. “The Agenda for Structural Reform in Europe.” Forthcoming in *Growth and Reform: European Economies in the Wake of the Economic Crisis*. Oxford University Press.

Fernald, John G. 2007. “Trend Breaks and Contractionary Technology Improvements.” *Journal of Monetary Economics.*

Fernald, John G. 2014. “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity.” Working Paper no. 2012-19. Federal Reserve Bank of San Francisco.

Fernald, John G. 2015. “Productivity and Potential Output Before, During, and After the Great Recession.” *NBER Macroeconomics Annual* 29: 1–51.

Fernald, John G. and J. Christina Wang. 2016. “Why Has the Cyclicality of Productivity Changed? What Does It Mean?” *Annual Review of Economics*. http://www.annualreviews.org/doi/full/10.1146/annurev-economics-080315-015018

Field, Alexander J. (2003). "The Most Technologically Progressive Decade of the Century." *American Economic Review* 93(4): 1399-1413.

Goldschlag, Nathan and Alexander T. Tabarrok. 2014. “Is Regulation to Blame for the Decline in American Entrepreneurship?” GMU Working Paper in Economics No: 15-11.

Gordon, Robert J. “The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections.” Working Paper no. 19895. National Bureau of Economic Research

Gordon, Robert J. 2016. “The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War.” Princeton University Press.

Gutiérrez, Germán and Thomas Philippon. 2016. “Investment-less Growth: An Empirical Investigation.” NBER Working Paper 22897.

Hall, Robert and Nicolas Petrosky-Nadeau. 2016. “Changes in Labor Participation and Household Income,” Economic Letter, Federal Reserve Bank of San Francisco,

Hall, Robert E. 2016. “Macroeconomics of Persistent Slumps.” Working Paper no. 22230, National Bureau of Economic Research.

Huang, Yu-Fan, Sui Luo, and Richard Startz 2016. “Are Recoveries All the Same: GDP and TFP?” Manuscript, University of California Santa Barbara.

Jones, Charles A. 2002. “Sources of U.S. Economic Growth in a World of Ideas.” *American Economic Review* 92 (1): 220–239.

Jorgenson, Dale W. and Zvi Griliches 1967. "The Explanation of Productivity Change." *The Review of Economic Studies* 34(3): 249–283.

Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh. 2008. "A Retrospective Look at the U.S. Productivity Growth Resurgence." *Journal of Economic Perspectives* 22(1): 3–24.

Kahn, James and Robert Rich 2007. "Tracking the New Economy: Using Growth Theory to Identify Changes in Trend Productivity." *Journal of Monetary Economics*:1670-1701.

Lansing, Kevin J. and Benjamin Pyle. 2015. "Persistent Overoptimism about Economic Growth." FRBSF Economic Letter, 2015-03.

Liu, Zheng, and Pengfei Wang. 2014. "Credit Constraints and Self-Fulfilling Business Cycles." *American Economic Journal: Macroeconomics* 6(1): 32–69.

Müller, Ulrich.K. and Mark W. Watson. 2008. "Testing Models of Low-Frequency Variability," *Econometrica* 76(5): 979–1016.

Nalewaik, Jeremy J. 2010. "The Income- and Expenditure-Side Measures of Output Growth." *Brookings Papers on Economic Activity* vol. 1: 71–106.

Nyblom, J. 1989. "Testing for the Constancy of Parameters Over Time. " *Journal of the American Statistical Association* 84: 223–230.

Okun, Arthur M. 1962. "Potential GNP: Its Measurement and Significance." *Proceedings* of the Business and Economics Statistics Section of the American Statistical Association: 98–104.

Oliner, Stephen D., Daniel E. Sichel, and Kevin Stiroh, 2007. "Explaining a Productive Decade." *Brookings Papers on Economic Activity*.

Oulton, Nicholas and María Sebastiá-Barriel 2013. "Effects of Financial Crises on Productivity, Capital and Employment." Paper Prepared for the IARIW-UNSW Conference on Productivity: Measurement, Drivers and Trends. Sydney, Australia, November 26-27, 2013.

Petrosky-Nadeu, Nicolas 2013. "TFP during a Credit Crunch." *Journal of Economic Theory* 148(3).

Polivka, Anne E. and Stephen M. Miller. "The CPS After the Redesign: Refocusing the Economic Lens." 1995. Bureau of Labor Statistics.

Reifschneider, Dave, William Wascher, and David Wilcox 2015. "Aggregate Supply in the United States: Recent Developments and Implications for the Conduct of Monetary Policy." *IMF Economic Review* vol. 63, no. 1: 71–109.

Robinson Peter M. 1988. "Root-N-consistent Semiparametric Regression." *Econometrica* 56: 931-954.

Sedlacek, Petr and Vincent Sterk, 2016. "The Growth Potential of Startups over the Business Cycle." Manuscript, University College London.

Solow, Robert M. 1957. "Technical Change and the Aggregate Production Function." *The Review of Economics and Statistics* 39 (3): 312–320.

Stock, James H. 1989. "Nonparametric Policy Analysis." *Journal of the American Statistical Association* 84(406): 567-575.

Stock, J.H., and M.W. Watson 1998. "Median Unbiased Estimation of Coefficient Variance in a Time Varying Parameter Model." *Journal of the American Statistical Association* 93: 349-358.

Stock, J.H. and M.W. Watson 2016. "Factor Models and Structural Vector Autoregressions in Macroeconomics." Forthcoming, *Handbook of Macroeconomics*, v. 2.

Syverson, Chad 2013. "Will History Repeat Itself? Comments on 'Is the Information Technology Revolution Over?'" *International Productivity Monitor*, vol. 25: 37-40.

Syverson, Chad. 2016. "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown." Working Paper no. 21974. Cambridge, Mass.: National Bureau of Economic Research.

Van Reenen, John, Nicholas Bloom, Mirko Draca, Tobias Kretschmer, Raffaella Sadun, Henry Overman, and Mark Schankerman 2010. "The Economic Impact of ICT." Research report, SMART N. 2007/0020.

Yellen, Janet L. 2016. "Macroeconomic Research After the Crisis." At "The Elusive 'Great' Recovery: Causes and Implications for Future Business Cycle Dynamics" 60th annual economic conference sponsored by the Federal Reserve Bank of Boston, Boston, Massachusetts, October 14.

Zhang, Ting and Wei Biao Wu. 2012. "Inference of Time-Varying Regression Models," *Annals of Statistics* 40(3): 1376-1402.

Figure 1: Unemployment and Output

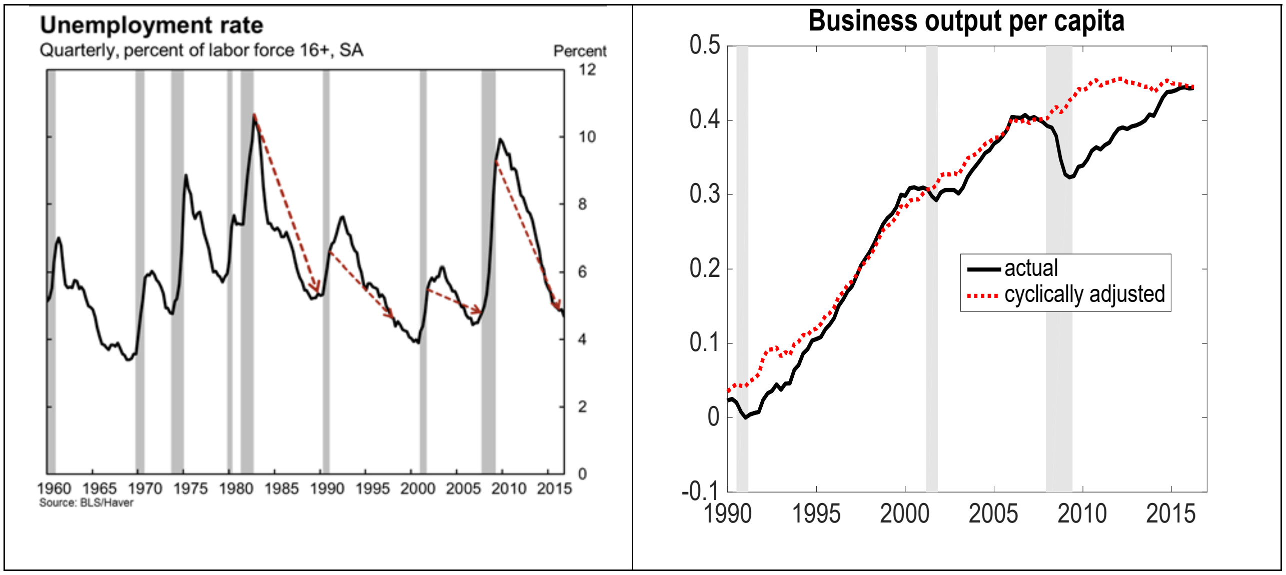

Notes: In the left panel, arrows connect the unemployment rate at the NBER-dated troughs with the rate 28 quarters later (or at the next peak, whichever comes first). In the right panel, the black line is the log of business output per person (normalized to 0 in 1991); the red line cyclically adjusts those data using Okun's Law as described in the text (normalized to equal the black line in 2007Q3).

Figure 2. SPF Forecasts of GDP and the Unemployment Rate, made in 2010 through 2015

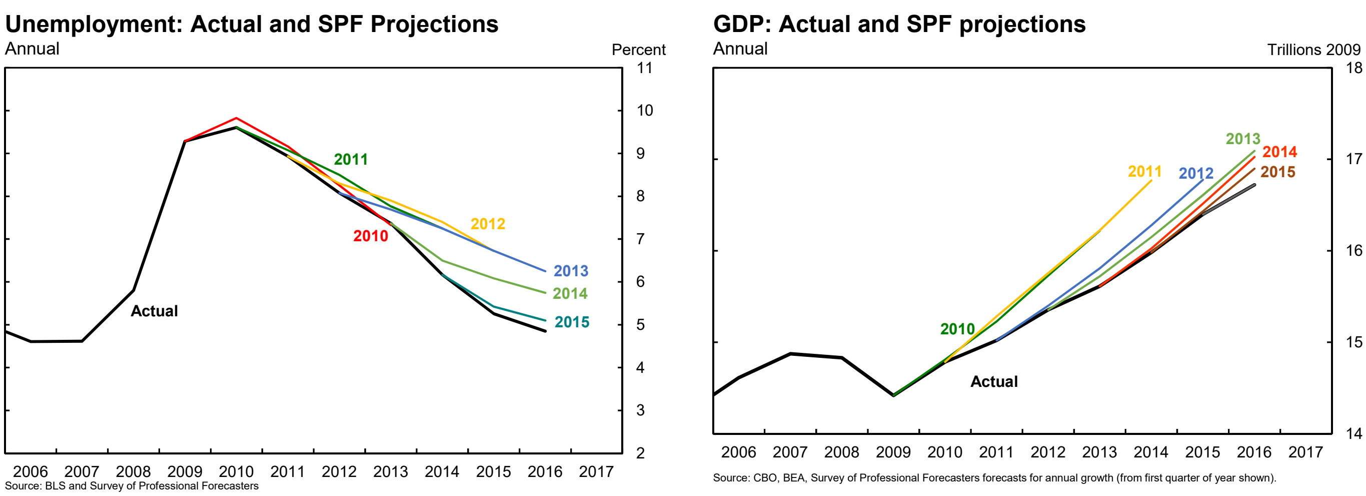

Notes: Median forecasts from the Survey of Professional Forecasters are from the first quarter of year indicated for annual averages of unemployment and GDP growth in that and subsequent years. The GDP figure on the right assumes the previous year's (revised) level is known and then projects using the published forecasts for annual growth rates. For example, the line for 2010 starts at 2009 actual, and uses 2010Q1 forecasts for annual growth in years 2010 on. The GDP figure follows Lansing and Pyle (2015).

Figure 3: Fraction of Employed People on Part Time for Economic Reasons, with Fitted Value from a Regression on Unemployment

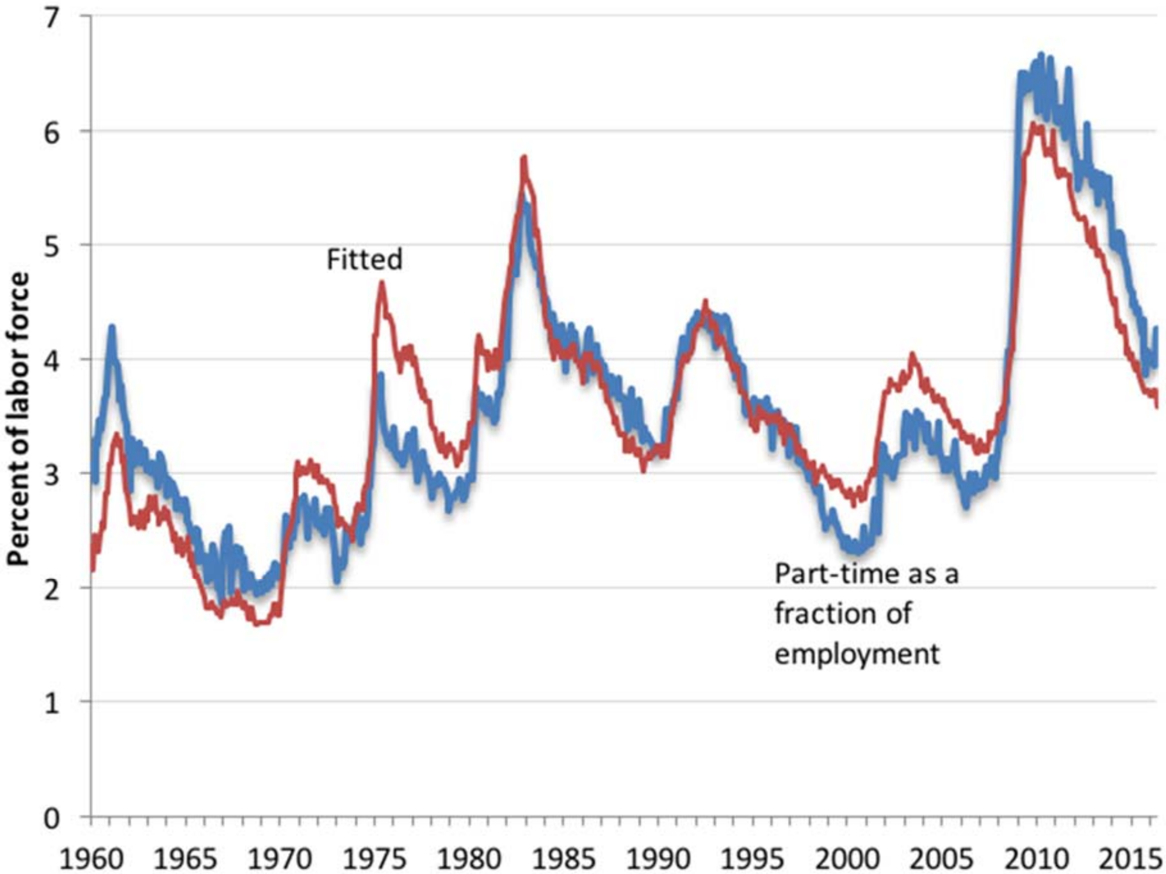

Figure 4: Data and Okun's Law Filtered Data: Output and Labor Productivity

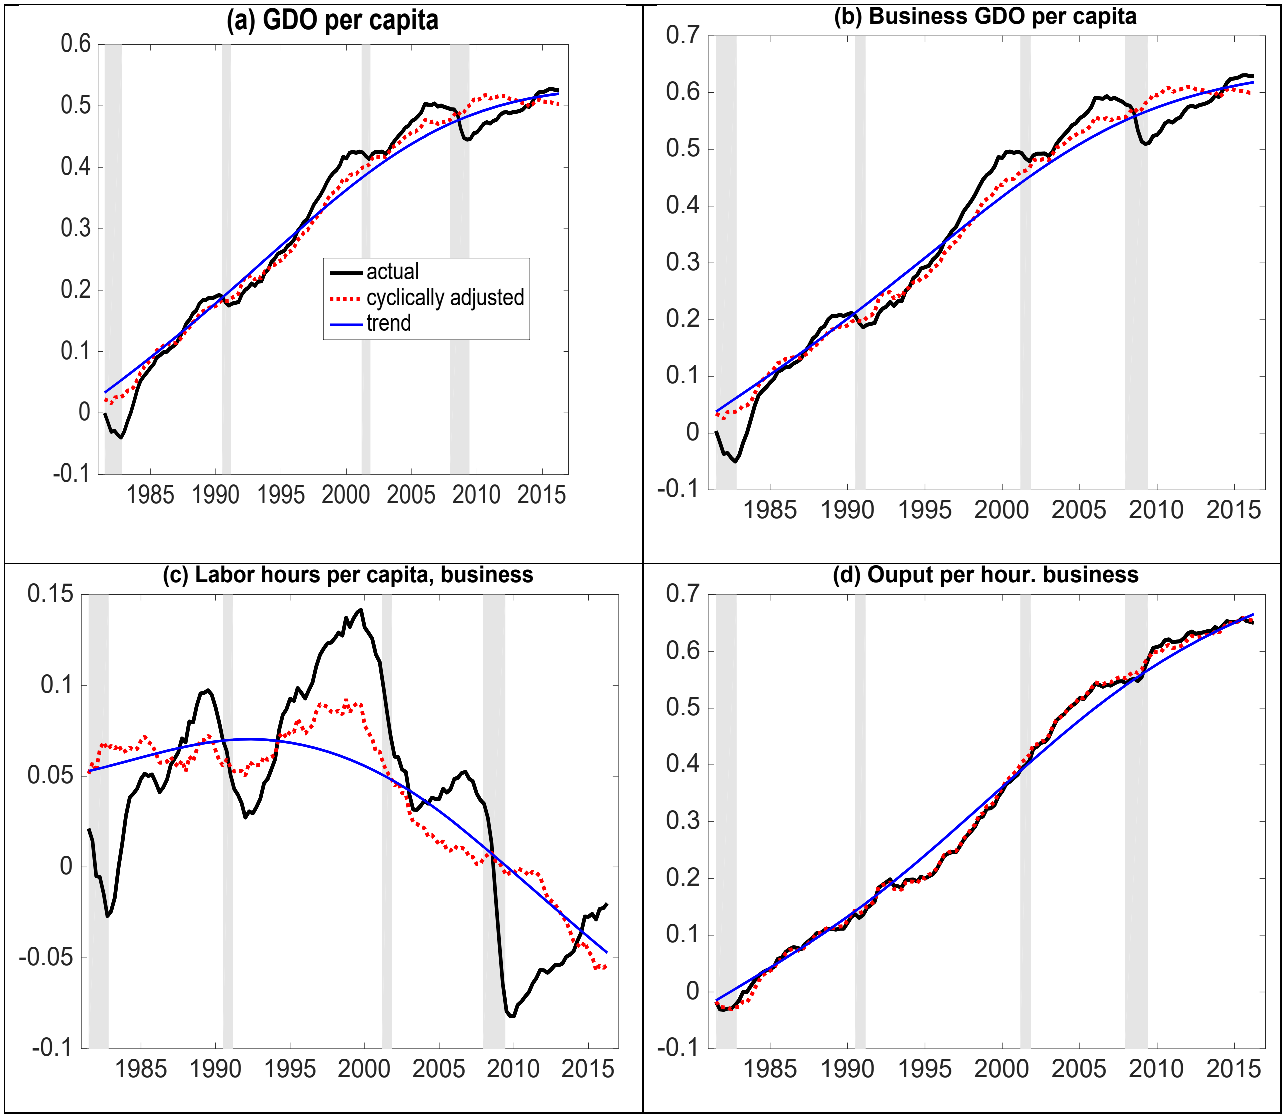

Notes: Plots of cumulated growth rates. Black lines are raw data, red lines are residuals (including constant terms) from Okun's Law regressions. Blue line is biweight filtered trend (bandwidth 60 quarters) fitted to the Okun's Law residuals. Levels are normalized to have the same means over the sample shown.

Figure 5: Data and Okun's Law Filtered data: Labor Market Variables

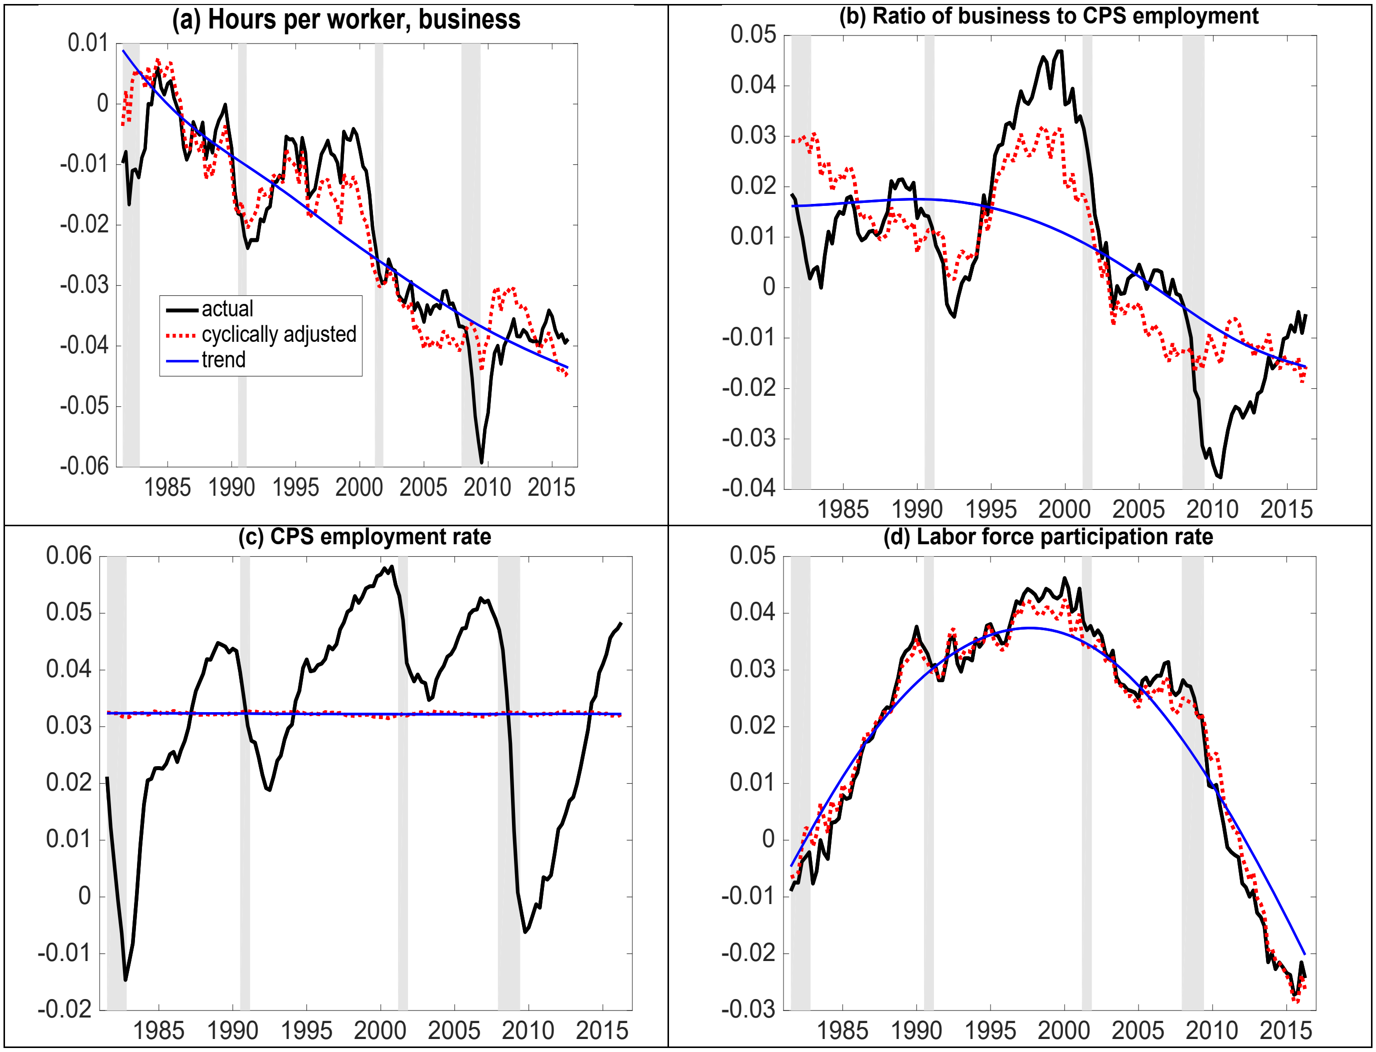

Notes: See Figure 4.

Figure 6: Data and Okun's Law Filtered Data:
Productivity, Capital Ratios, and Labor quality

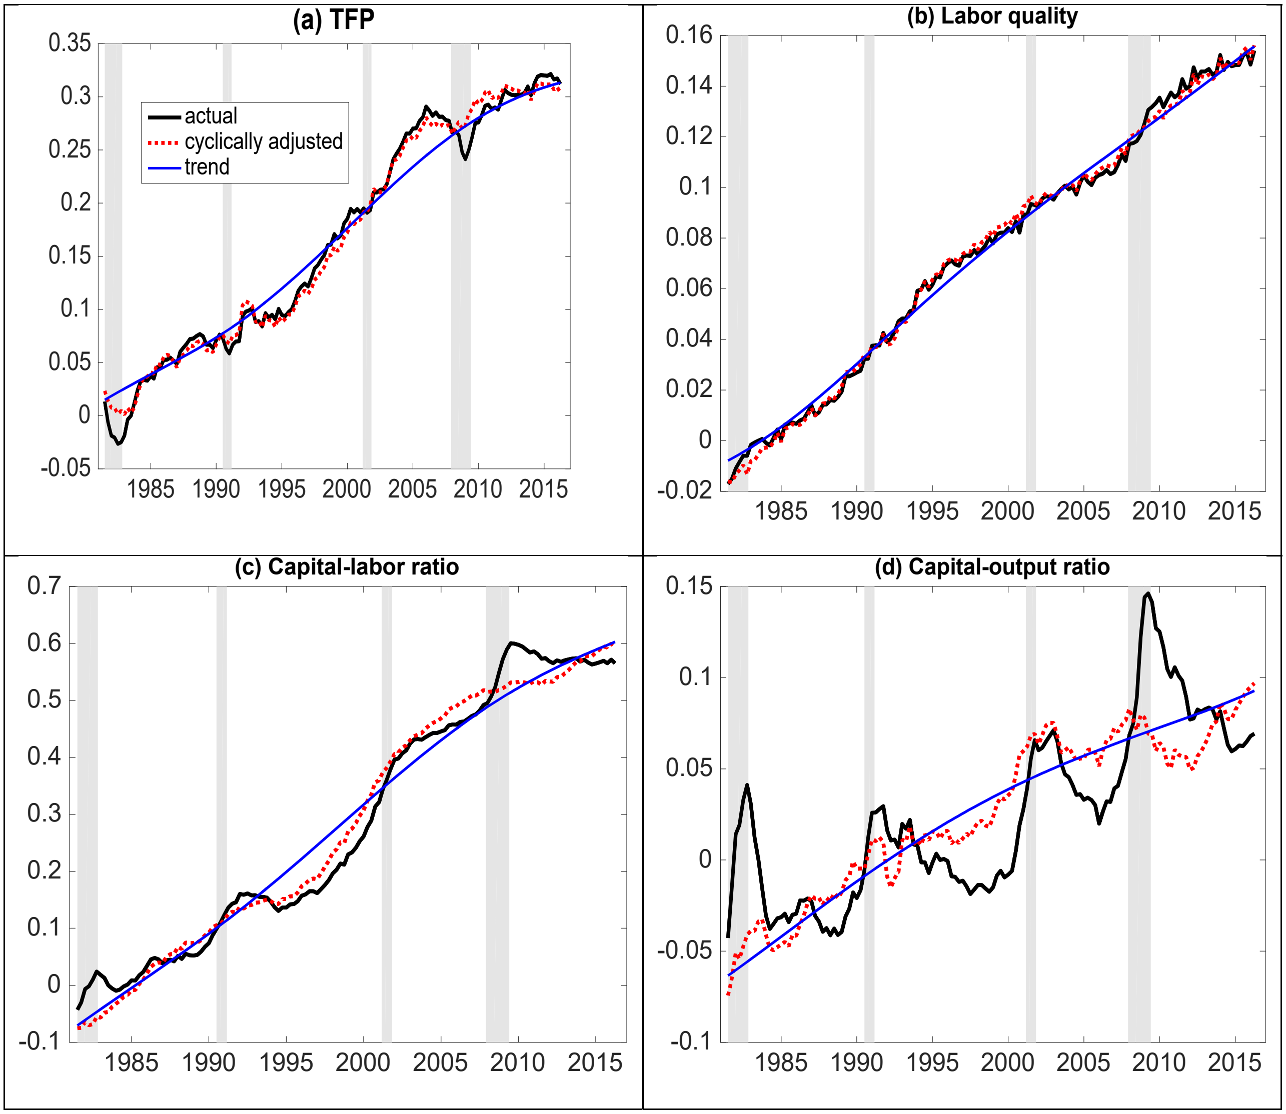

Notes: See Figure 4.

Figure 7: Forecasted and Actual Paths from the Factor Model: Growth Accounting Variables

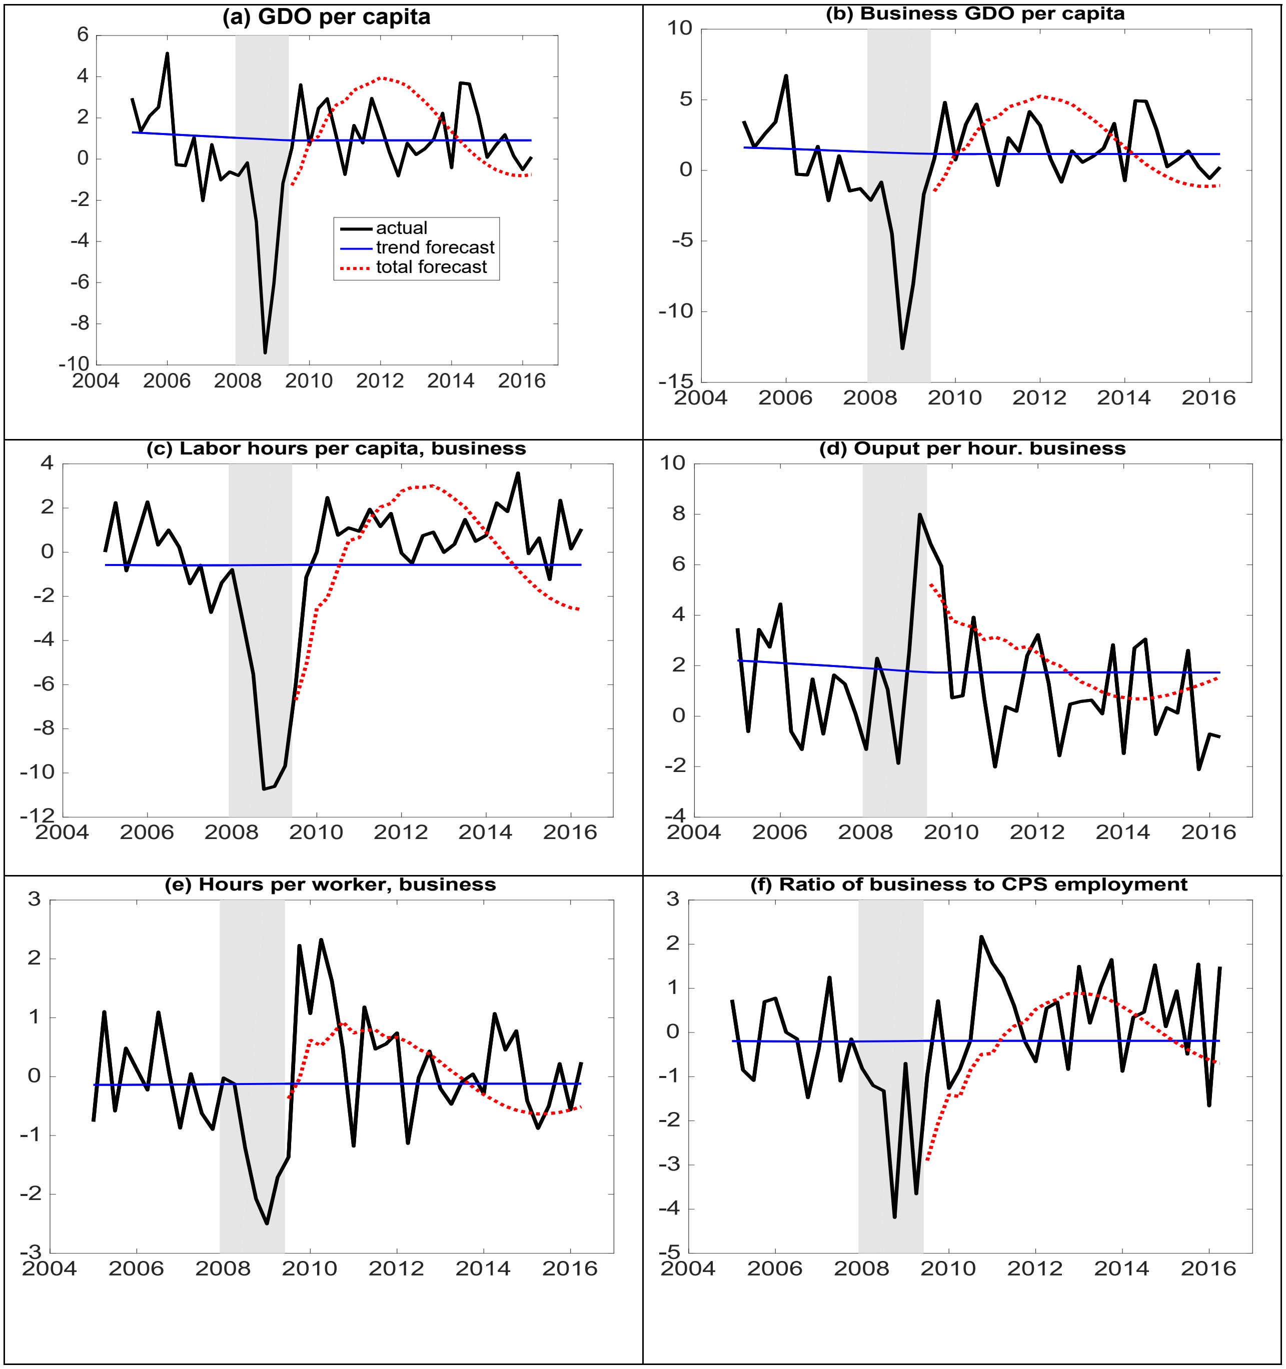

Figure 7, continued

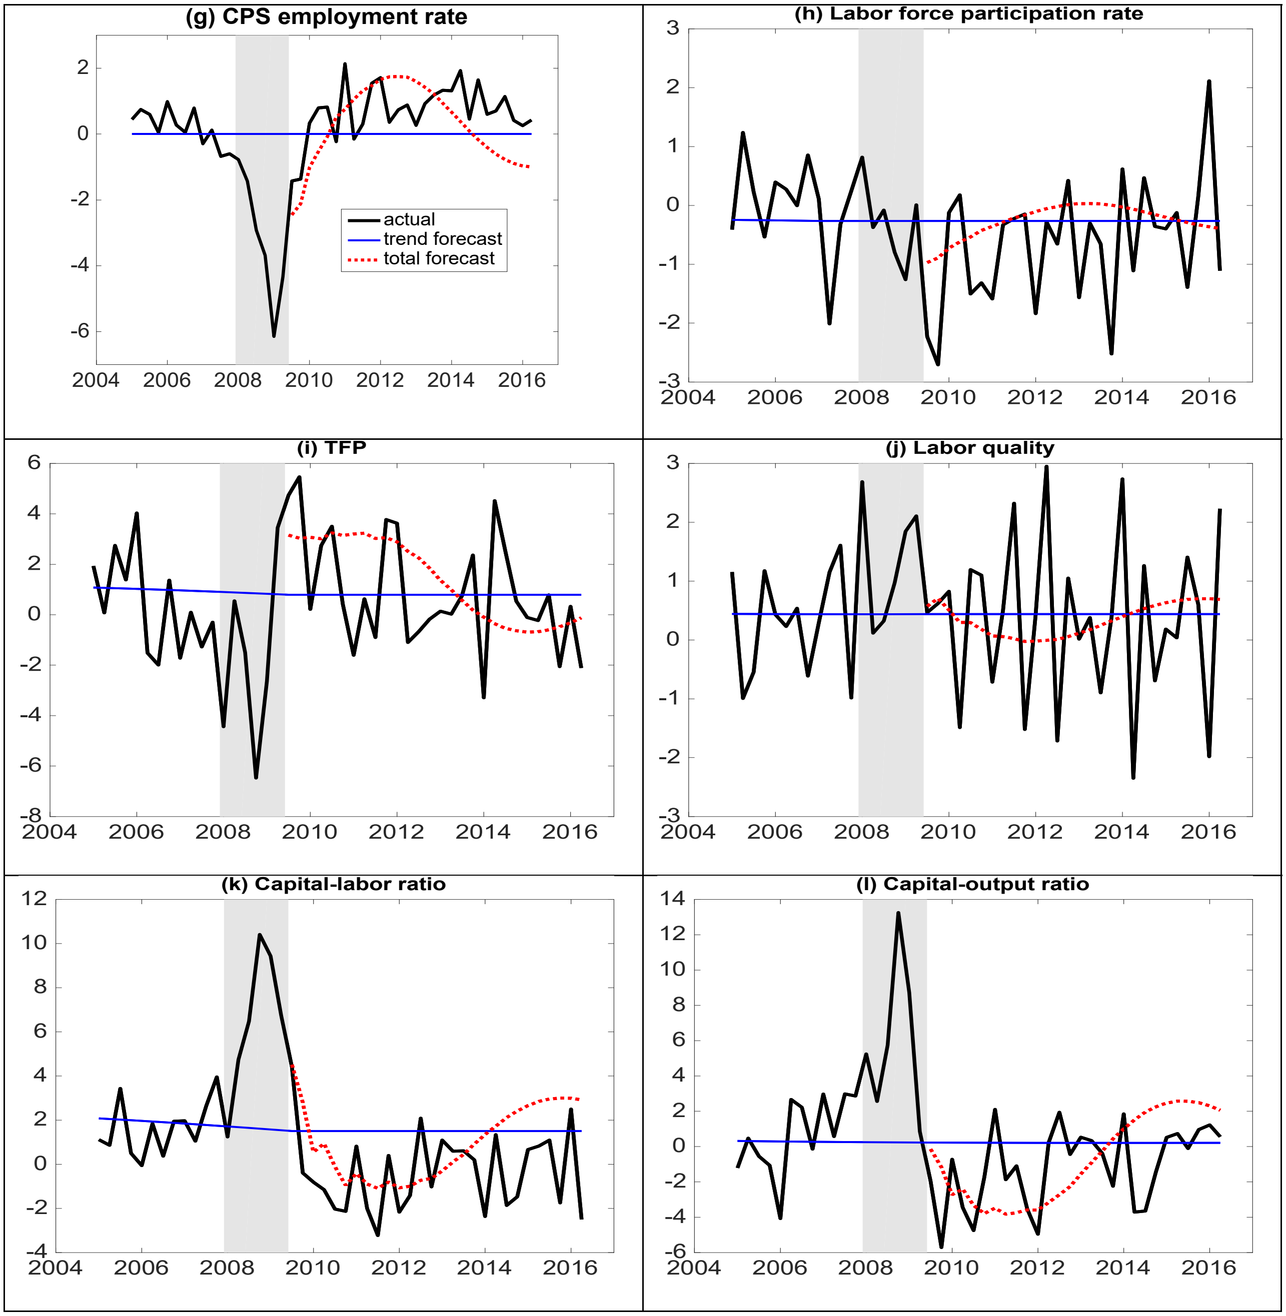

Notes to Figure 7: Black line is the actual growth rate of the variable, red line is its forecast based on the 6 factors, and the blue line is the long-term growth trend.

Figure 8. Forecasted and Actual Paths from the Factor Model: Selected NIPA Variables

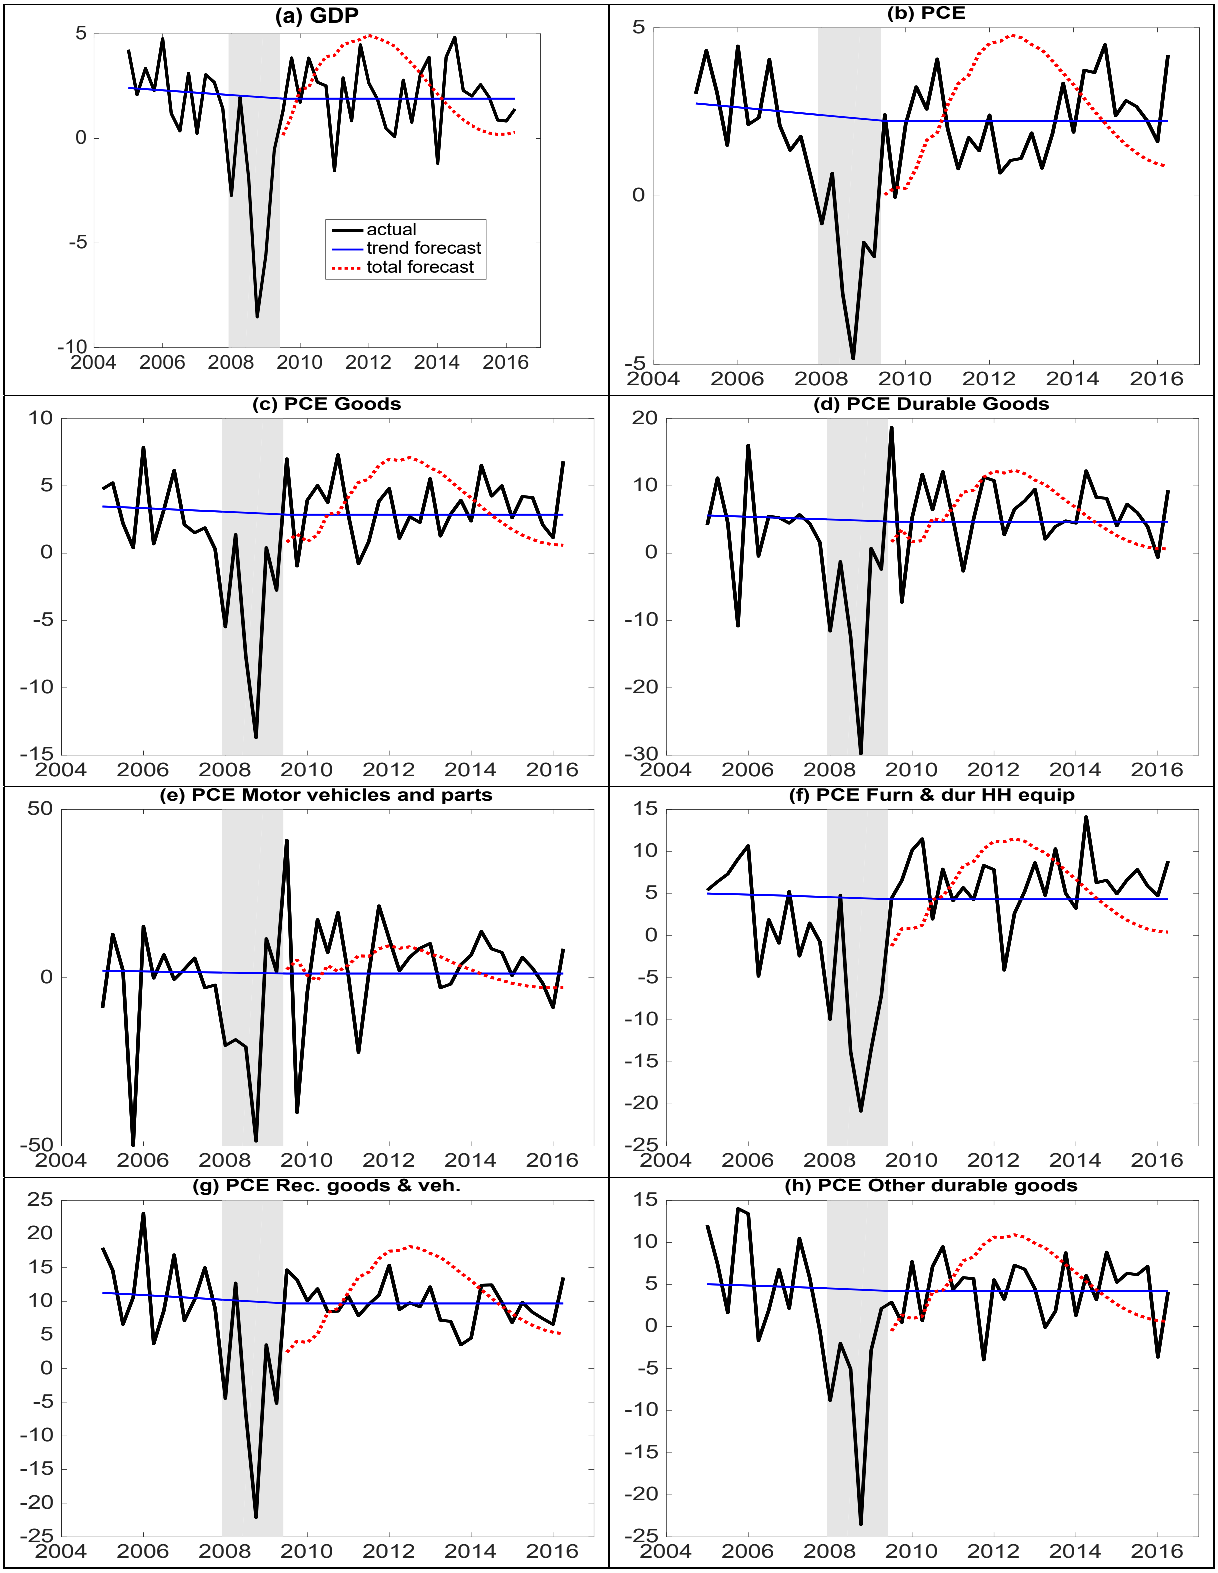

Figure 8, continued

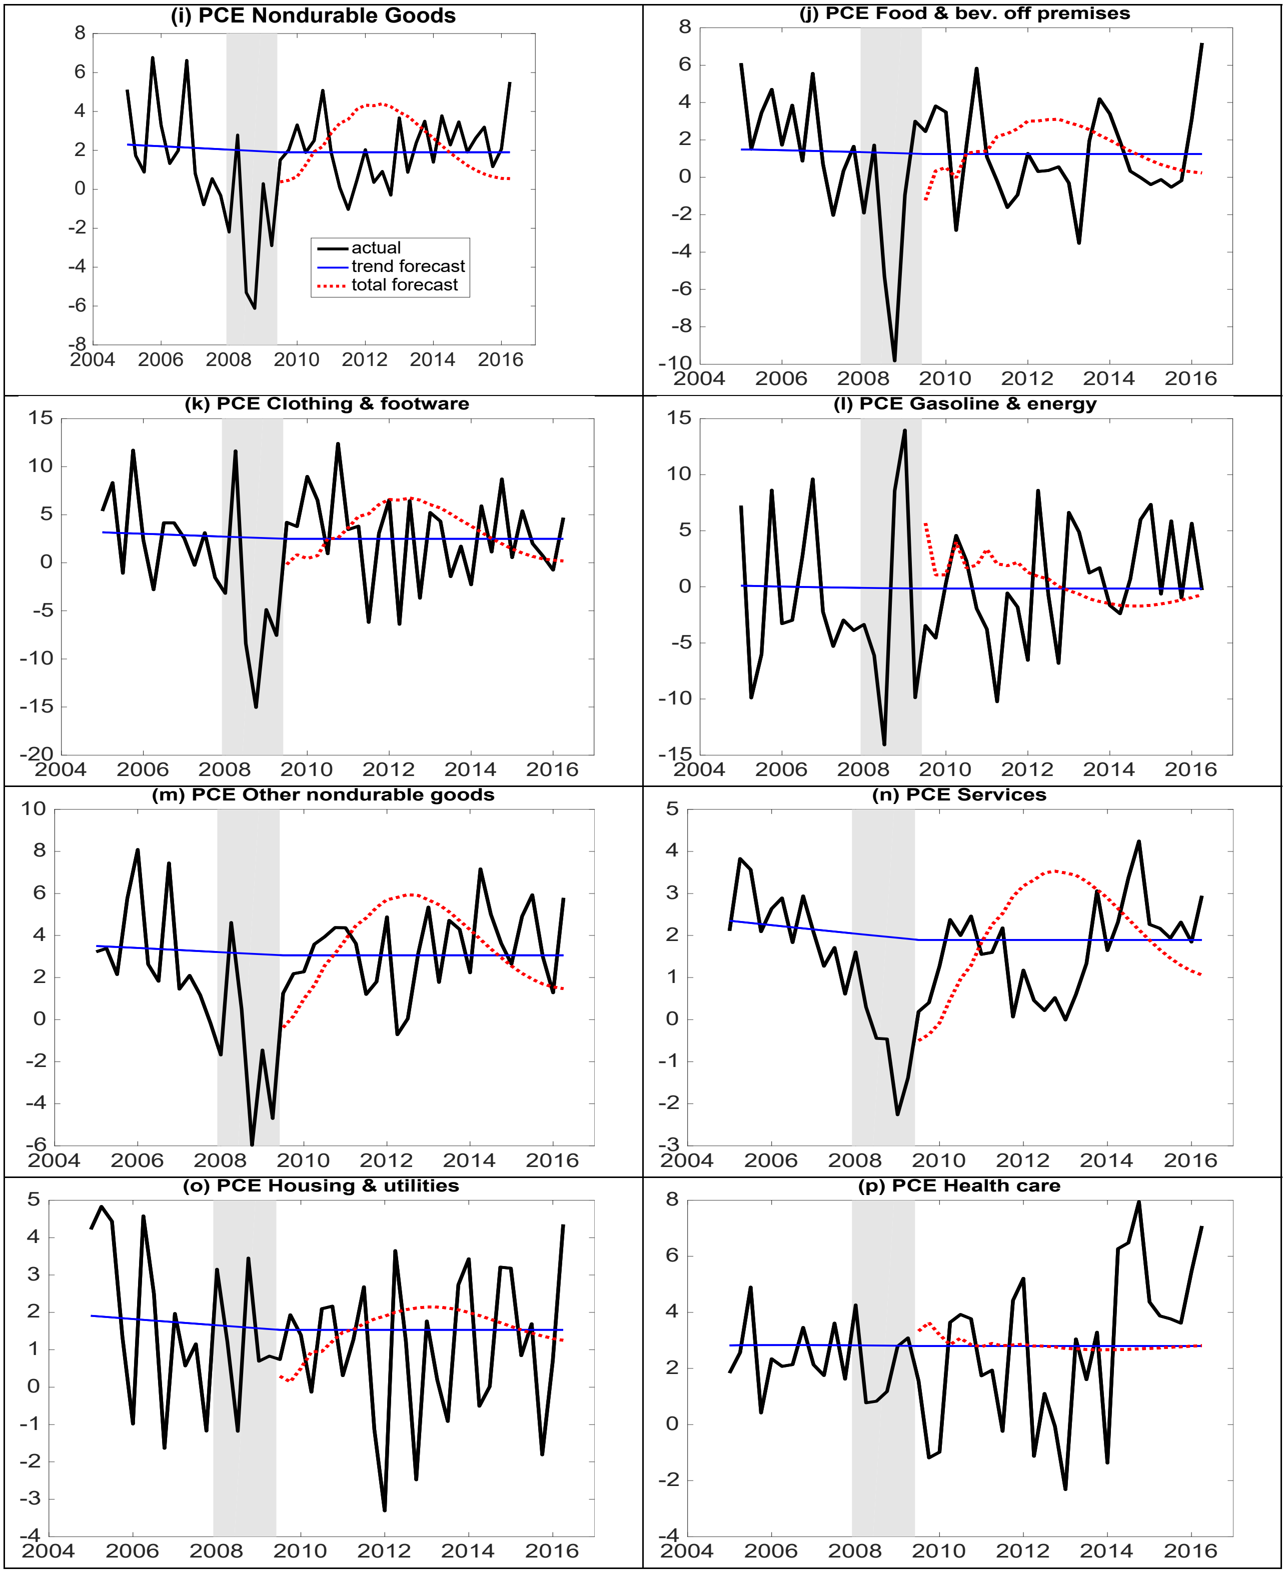

Figure 8, continued

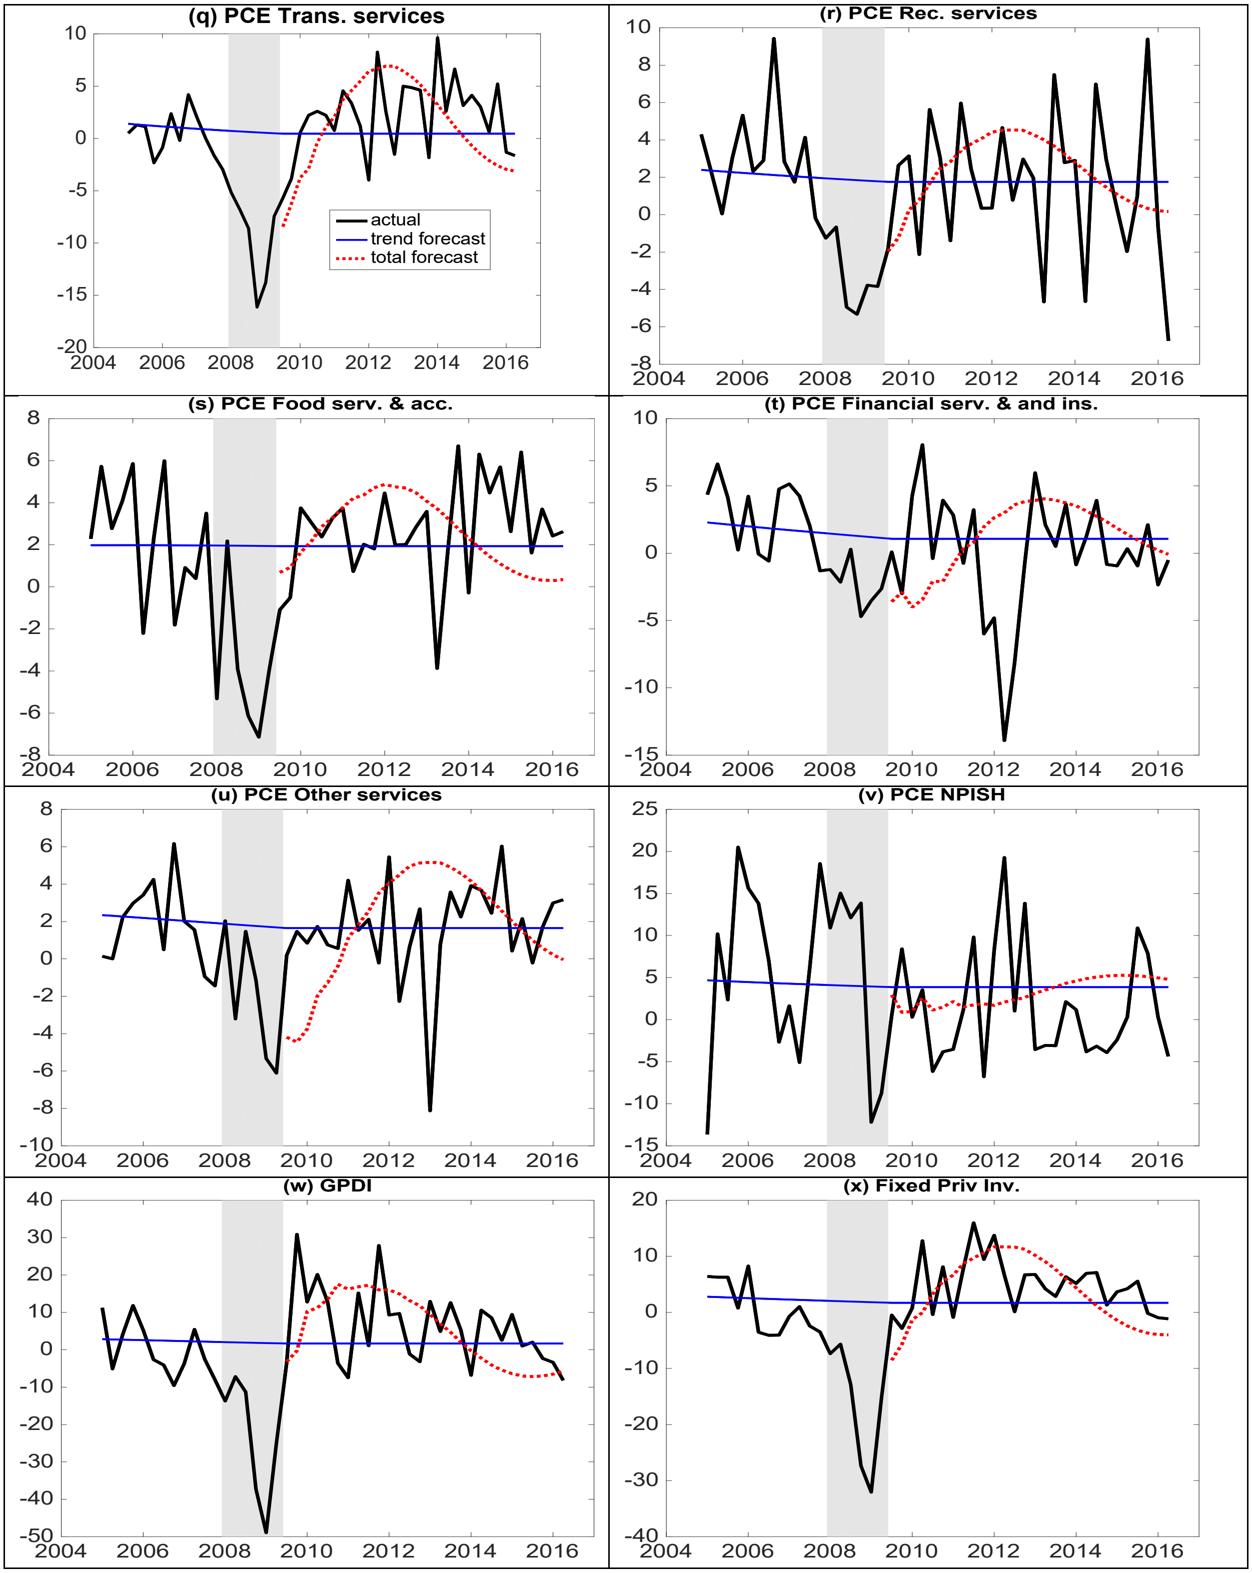

Figure 8, continued

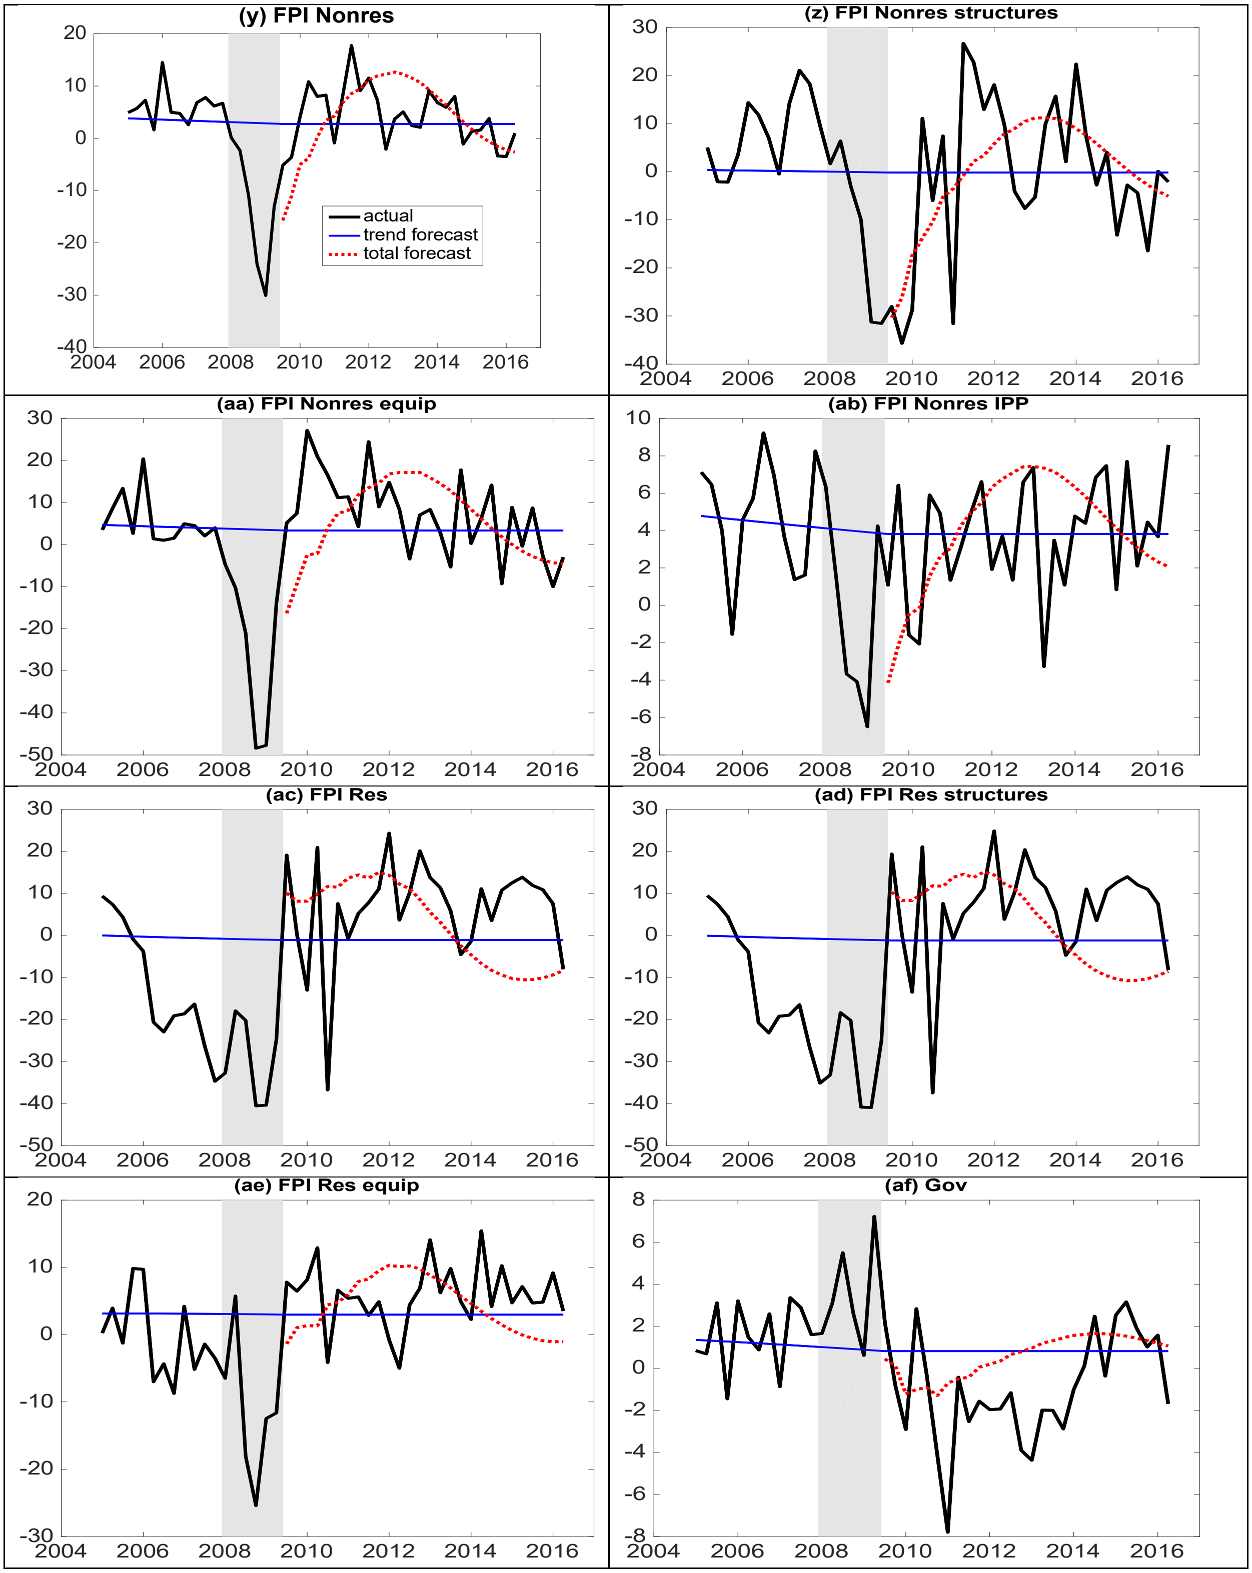

Figure 8, continued

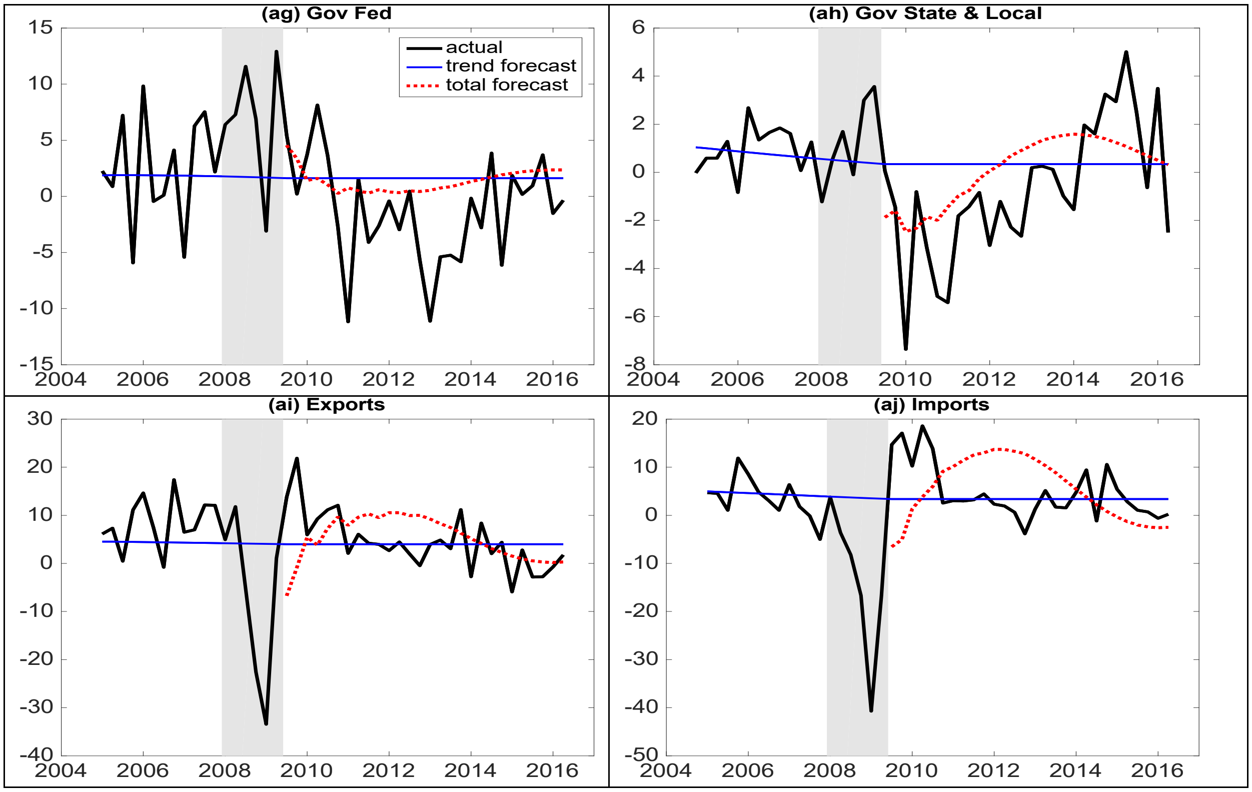

See the notes to Figure 7.

Figure 9. Forecasted and Actual Paths from the Factor Model: Employment variables

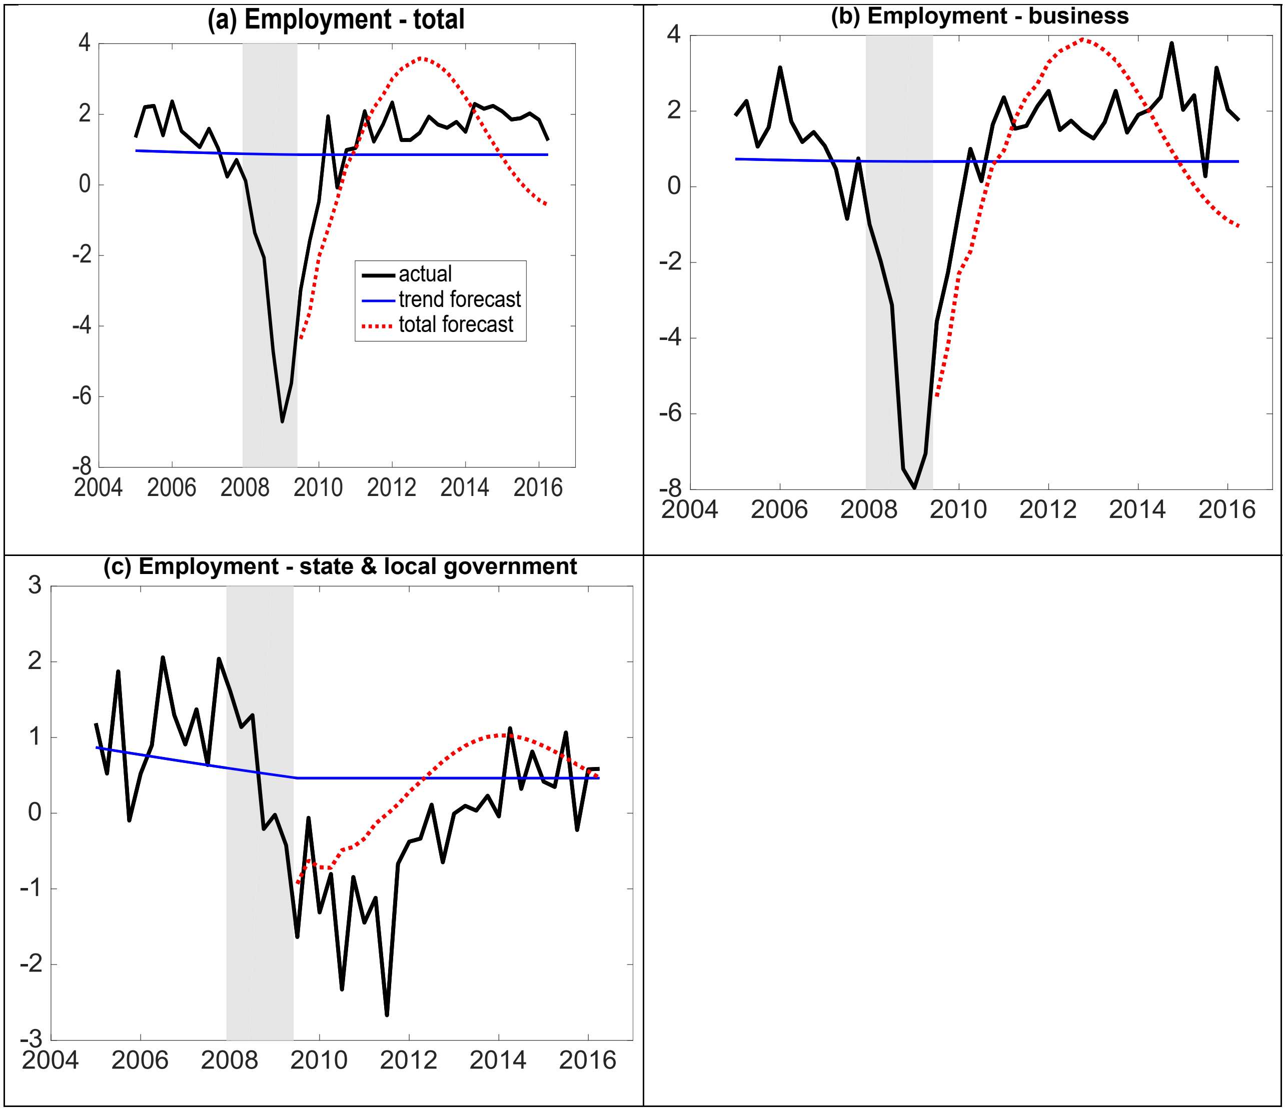

See the notes to Figure 7.

Figure 10: Real-Time Estimates of Prospective 10-year Growth in Labor Productivity

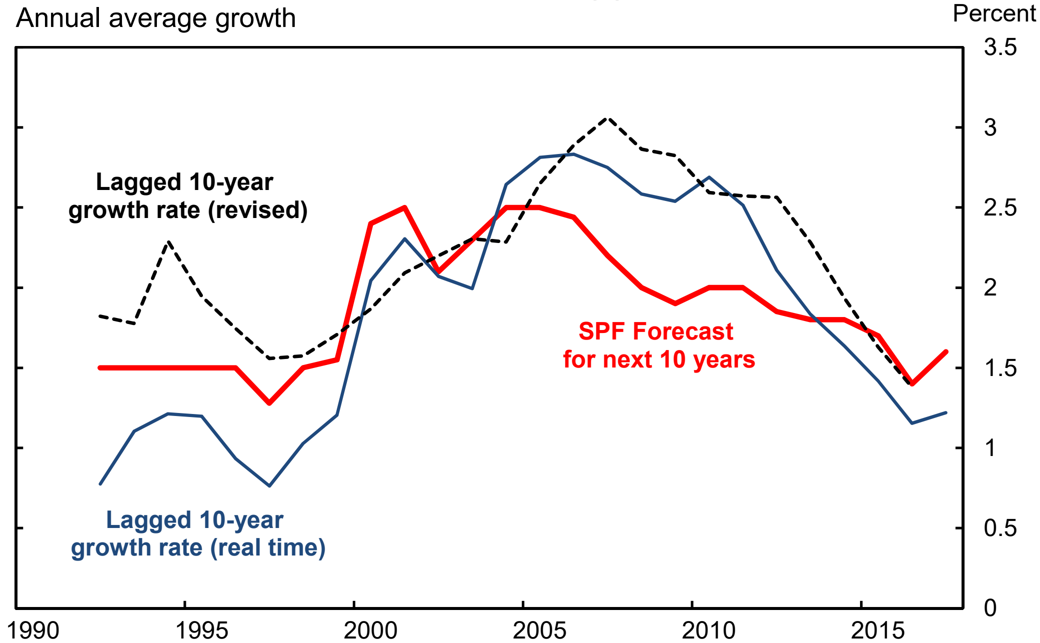

Source: Survey of Professional Forecasters, ALFRED (FRB St. Louis), and BLS. Revised actual data are from Feb. 2, 2017.. Output per hour is for the non-farm business sector. Surveys are from the first quarter of the year, and are the annual average over the next 10 years.

Figure 11: Cyclically-Adjusted TFP and Estimated Low-Frequency Mean Growth Rates

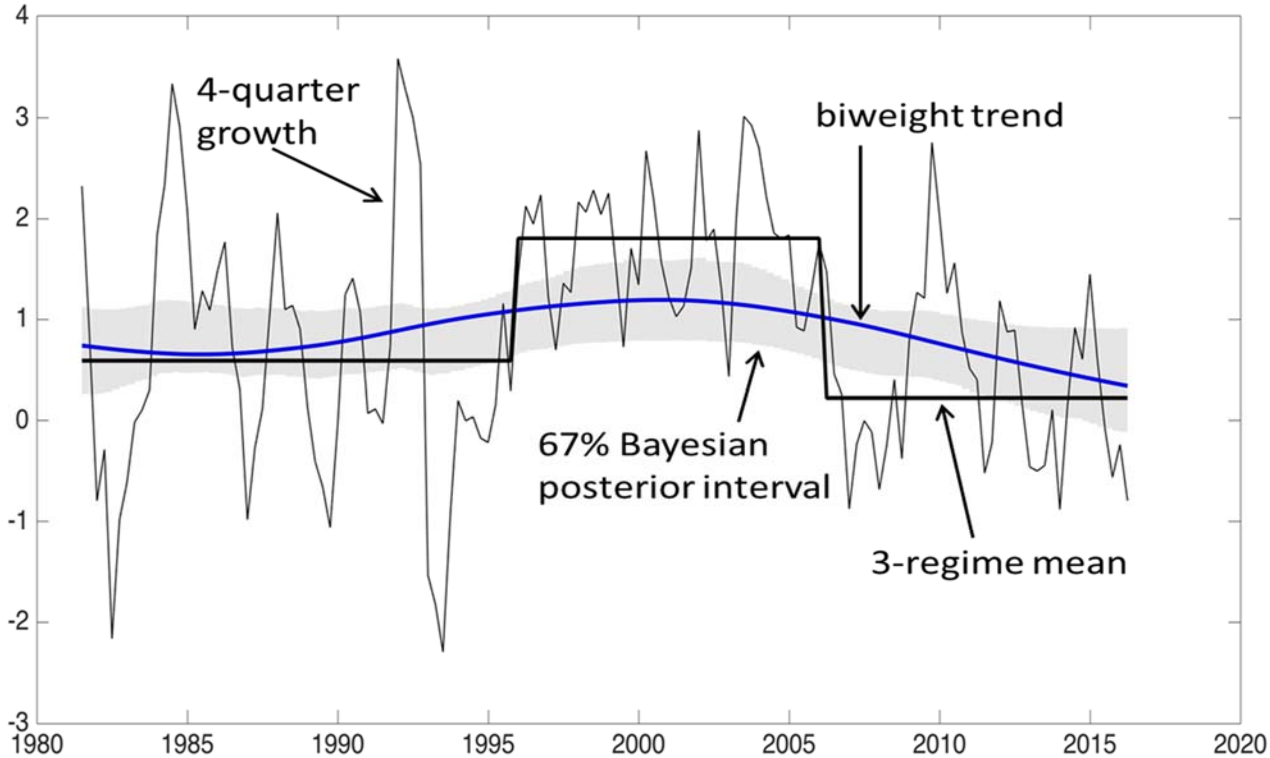

Notes: TFP is cyclically adjusted. The thin black line is its four-quarter growth rate. The blue line is the cyclically-adjusted trend using a biweight filter (60-quarter bandwidth). The shaded area is a 67% Bayes posterior set. The dark black line are the means estimated within the three regimes estimated by break tests, with break dates in 1995Q4 and 2006Q1 from Table 6.

Figure 12: Posterior Density of Date of Maximum Trend Growth in TFP, 1981-2016

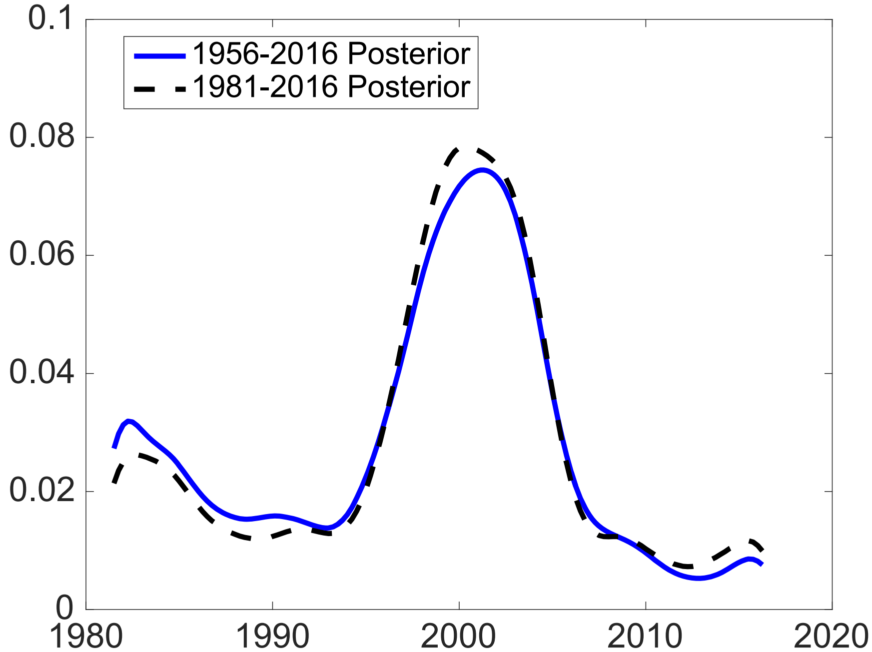

Notes: TFP is cyclically adjusted. Computed using Bayes implementation of the random walk-plus-noise model for productivity growth, as discussed in the text.

Figure 13. Investment in Productivity Improvements

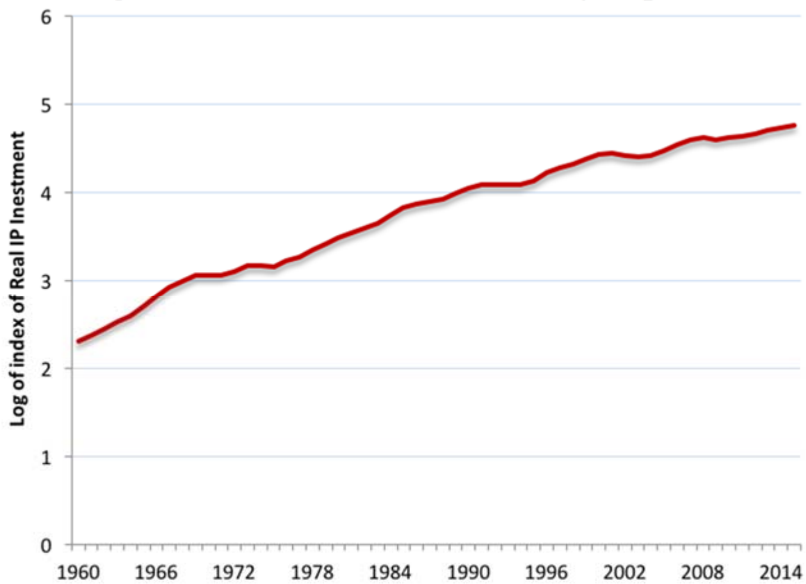

Figure 14. Equipment Investment

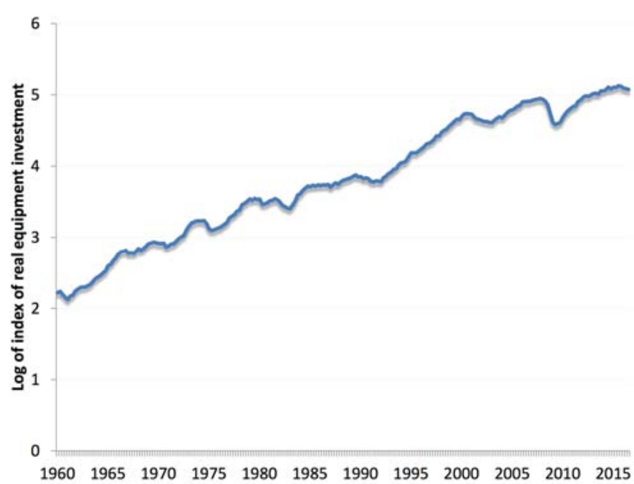

Figure 15. Business Earnings as a Ratio to the Value of Capital

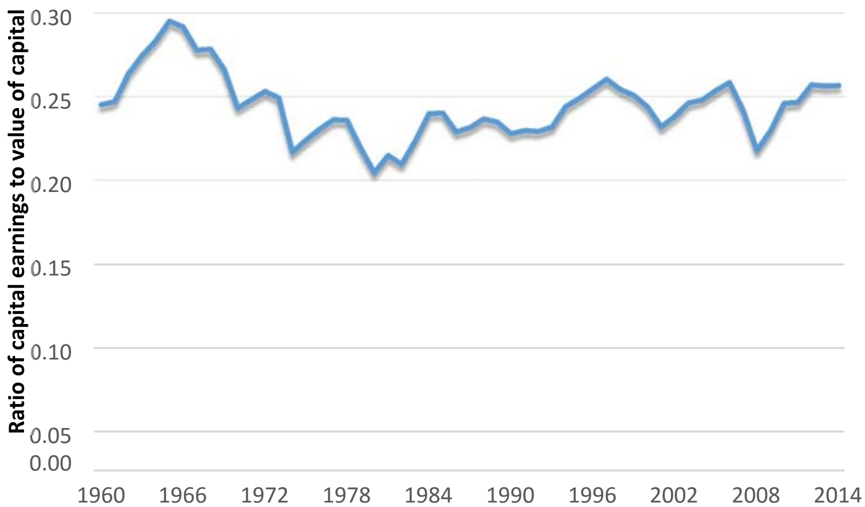

Figure 16: Labor-Force Participation Rates by Sex

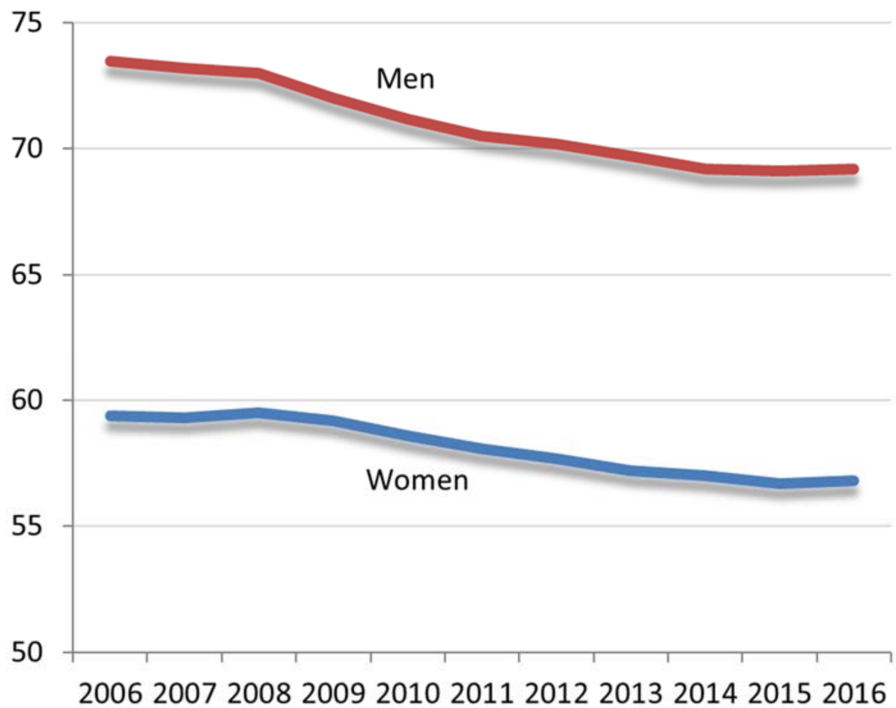

Figure 17. Labor-Force Participation Rate, Actual and Adjusted for Changing Demography

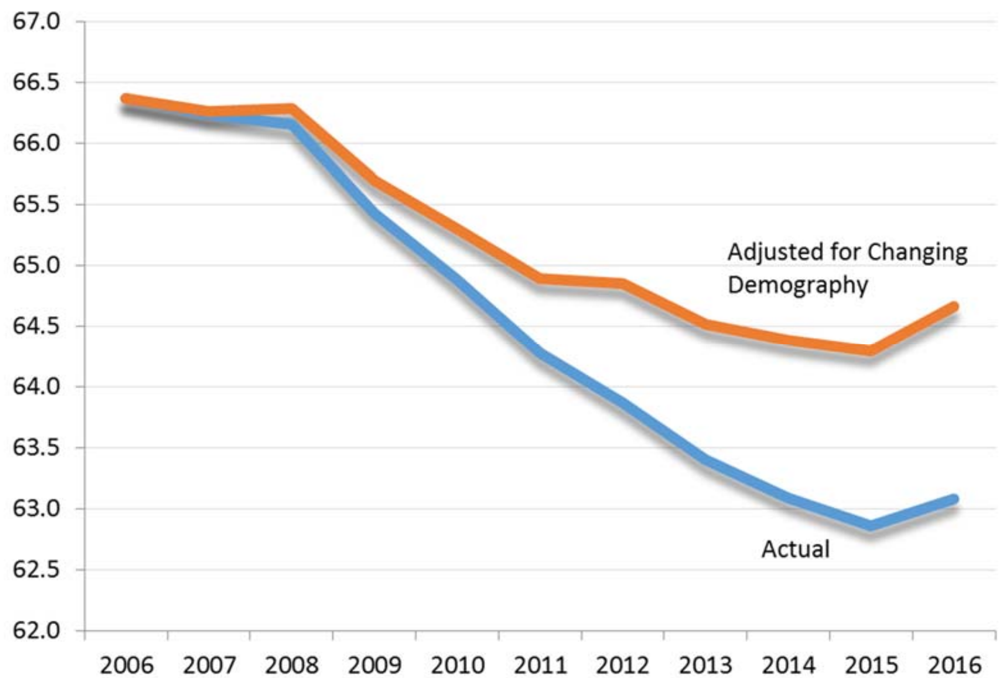

Figure 18. Role of Family Income in Participation Rates

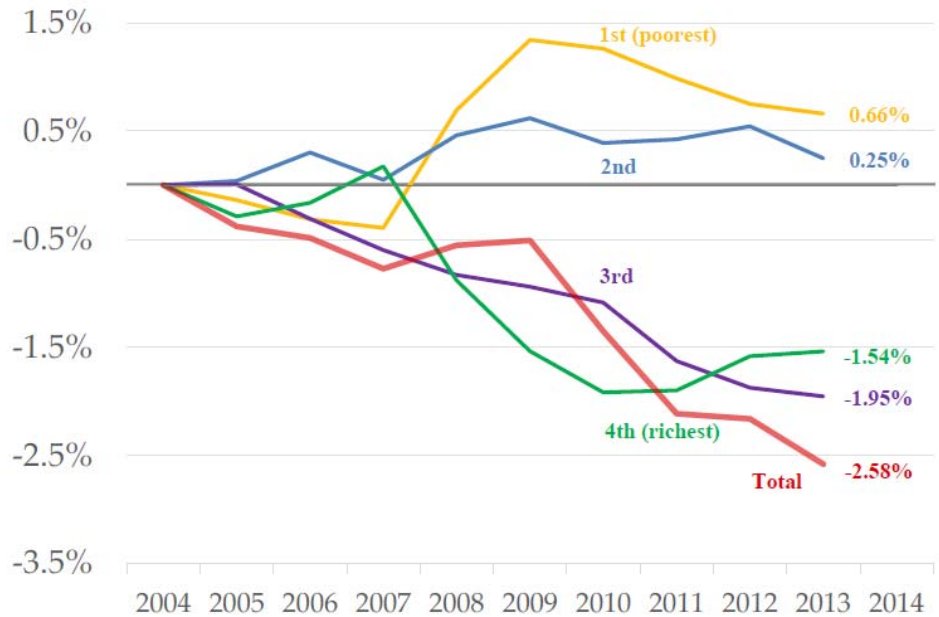

Table 1. Categories of Quarterly Time Series Used to Estimate the Factors

	Category	Number of series
(1)	NIPA	12
(2)	Industrial Production	7
(3)	Employment and Unemployment	30
(4)	Orders, Inventories, and Sales	8
(5)	Housing Starts and Permits	6
(6)	Prices	24
(7)	Productivity and Labor Earnings	5
(8)	Interest Rates	9
(9)	Money and Credit	5
(10)	International	9
(11)	Asset Prices, Wealth, Household Balance Sheets	9
(12)	Oil Market Variables	6
	Total	123

Notes: For the full list of series and data transformations see the supplemental data appendix.

Table 2. Cyclicality of Real Output and its Components

	Generalized Okun's law coefficient and std. error	Standard deviations of components			R^2 from regressing on factors
		cycle (c)	trend(μ)	irregular (z)	
(1) GDP	-1.49 (0.18)	1.90	0.58	1.77	0.66
(2) GDO (Average of GDP, GDI)	-1.53 (0.17)	1.92	0.57	1.61	0.72
(3) Business GDO	-2.03 (0.21)	2.53	0.59	2.11	0.73
(4) GDP per capita	-1.48 (0.17)	1.88	0.52	1.84	0.60
(5) GDO per capita	-1.52 (0.17)	1.89	0.51	1.63	0.67
(6) Business GDO per capita	-2.02 (0.20)	2.51	0.54	2.12	0.70
(7) Total factor productivity	-0.50 (0.19)	1.24	0.24	2.27	0.38
(8) α*Capital/Pop.	-0.09 (0.06)	0.20	0.19	0.32	0.37
(9) (1-α)*(Lab Qual * Hours/Pop.)	-1.43 (0.14)	1.54	0.26	1.24	0.57
(10) Bus. labor hours per capita	-2.30 (0.19)	2.54	0.36	1.51	0.74
(11) Hours per worker, business	-0.35 (0.1)	0.55	0.04	1.05	0.25
(12) Ratio of bus.empl to CPS empl	-0.71 (0.09)	0.73	0.08	1.20	0.24
(13) CPS employment rate	-1.08 (0.01)	1.36	0.00	0.10	0.89
(14) Labor-force participation rate	-0.16 (0.10)	0.32	0.33	0.87	0.02
(15) Bus. output per hour (labor prod.)	0.28 (0.22)	0.77	0.37	2.23	0.24
(16) TFP / (1 -α)	-0.75 (0.29)	1.88	0.35	3.41	0.39
(17) Capital-Output ratio × α/(1-α)	0.90 (0.09)	1.30	0.07	1.09	0.75
(18) Labor quality	0.13 (0.05)	0.37	0.05	0.99	0.06

Notes: The Okun's law coefficients are $\beta(1)/4$, so they are measured in quarterly percentage points of growth per percentage point change in the unemployment rate. The standard deviations of the components are for quarterly growth rates reported in percentage points at an annual rate. The R^2 is from the regression of the variable on the factors used in factor model.

Table 3: Shortfall of the Post-Crisis Recovery Relative to Earlier Recoveries: Growth Accounting Decomposition Using Okun's Law Cyclical Adjustment

	Historical values (not cyclically adjusted)			**Cyclically adjusted**					
						Annual shortfall			
	Three previous recovs.	2009Q2-2016Q2	Annual shortfall (a)-(b)	Three previous recovs.	2009Q2-2016Q2	**Cyclically adjusted shortfall (d) - (e)**	*Shortfall in smooth trend*	*Residual shortfall (f) - (g)*	Cumul. shortfall
	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)
(1) GDP	3.60	2.06	1.54	2.95	0.96	1.99	1.26	0.73	14.94
(2) GDO (Average of GDP, GDI)	3.57	2.20	1.37	2.92	1.11	1.81	1.24	0.57	13.54
(3) Business GDO	4.04	2.76	1.29	3.18	1.29	1.89	1.31	0.58	14.14
(4) GDP per capita	2.48	1.02	1.45	1.84	-0.07	1.91	1.13	0.78	14.30
(5) GDO per capita	2.45	1.16	1.29	1.80	0.07	1.73	1.11	0.62	12.90
(6) Business GDO per capita	2.92	1.72	1.21	2.07	0.26	1.81	1.18	0.63	13.49
(7) Total factor productivity	1.30	0.89	0.42	0.99	0.28	0.71	0.36	0.35	5.12
(8) α*Capital/Pop.	0.79	0.24	0.55	0.77	0.24	0.53	0.40	0.13	3.78
(9) (1-α)*(Lab Qual * Hours/Pop.)	0.83	0.59	0.24	0.30	-0.27	0.57	0.41	0.15	4.04
(10) Bus. labor hours per capita	0.81	0.63	0.18	-0.06	-0.76	0.70	0.55	0.14	5.00
(11) Hours per worker, business	0.07	0.24	-0.17	-0.10	-0.07	-0.03	-0.06	0.03	-0.24
(12) Ratio of bus.empl to CPS empl	0.12	0.37	-0.25	-0.11	0.01	-0.12	0.06	-0.18	-0.83
(13) CPS employment rate	0.43	0.68	-0.25	0.00	0.00	0.00	0.00	0.00	0.02
(14) Labor-force participation rate	0.19	-0.66	0.85	0.15	-0.69	0.85	0.56	0.29	6.11
(15) Bus. output per hour (labor prod.)	2.11	1.09	1.03	2.12	1.01	1.11	0.62	0.49	8.09
(16) TFP / (1 -α)	1.95	1.44	0.51	1.48	0.51	0.96	0.48	0.48	6.98
(17) Capital-Output ratio × α/(1-α)	-0.26	-0.69	0.42	0.16	0.07	0.08	0.11	-0.02	0.59
(18) Labor quality	0.43	0.33	0.09	0.49	0.43	0.06	0.03	0.03	0.44

Notes: Entries are average annual percent changes or percentage point differences. Indented rows sum to next level of aggregation. Post-crisis recovery period is 2009Q2 through 2016Q2 (28 quarters). The three previous recoveries are the averages during the first 28 quarters from the troughs of 1982 and 1991, and the 24 quarters of the expansion after the 2001 trough. Cyclically-adjusted entries in columns (d) and (e) are residuals from Okun's Law regressions.

Table 4. Shortfall of the Post-Crisis Recovery Relative to 2009IV Forecasts: Growth Accounting Decomposition Using Forecast-Based Cyclical Adjustment

	Forecast	*Actual*	*Shortfall*	*(std. error)*
(1) GDP	2.63	2.06	0.57	(0.07)
(2) GDO (Average of GDP, GDI)	2.63	2.20	0.43	(0.07)
(3) Business GDO	3.11	2.76	0.35	(0.08)
(4) GDP per capita	1.51	1.02	0.48	(0.09)
(5) GDO per capita	1.51	1.16	0.35	(0.07)
(6) Business GDO per capita	1.99	1.72	0.27	(0.09)
(7) Total factor productivity	1.40	0.89	0.52	(0.09)
(8) α*Capital/Pop.	0.43	0.24	0.19	(0.01)
(9) (1-α)*(Lab Qual * Hours/Pop.)	0.15	0.59	-0.44	(0.05)
(10) Bus. labor hours per capita	-0.08	0.63	-0.72	(0.06)
(11) Hours per worker, business	0.08	0.24	-0.16	(0.03)
(12) Ratio of bus.empl to CPS empl	-0.16	0.37	-0.53	(0.06)
(13) CPS employment rate	0.26	0.68	-0.42	(0.02)
(14) Labor-force participation rate	-0.27	-0.66	0.40	(0.03)
(15) Bus. output per hour (labor prod.)	2.07	1.09	0.98	(0.08)
(16) TFP / (1 -α)	2.15	1.44	0.72	(0.12)
(17) Capital-Output ratio × α/(1-α)	-0.43	-0.69	0.26	(0.03)
(18) Labor quality	0.34	0.33	0.01	(0.04)

Notes: The first two numerical columns are forecasted and actual values of the variable in the first column, where the forecasts are computed using the factor model and the values of the factors through 2009q2. The third column is the shortfall (the negative of the forecast error), and the final column gives the standard error of the shortfall arising solely from sampling error in the estimated model parameters.

Table 5. Expected and Unexpected Contributions to GDP growth: NIPA Demand Components

	Growth Rate, 2009Q2-2016Q2	Average Share	Okun's Law Cyclical Adjustment						DFM Forecast		
			Okuns law coefficient (SE)	CA Growth Rate							
				Three previous recoveries	Post-crisis recovery	Total Shortfall	Trend shortfall	Irregular (z) shortfall	Forecast	Shortfall	SE
Real gross domestic product	2.06	1	-1.49 (0.18)	2.95	0.96	1.99	1.26	0.73	2.63	0.57	0.07
Personal consump. Expend.	1.54	0.68	-0.74 (0.14)	2.00	1.04	0.96	0.70	0.26	1.80	0.26	0.04
Goods	0.78	0.23	-0.44 (0.08)	0.80	0.48	0.32	0.24	0.08	0.86	0.08	0.03
Goods, durable	0.47	0.07	-0.25 (0.06)	0.43	0.28	0.15	0.12	0.03	0.50	0.03	0.03
Motor vehicles & parts	0.11	0.02	-0.09 (0.04)	0.10	0.04	0.06	0.09	-0.03	0.10	-0.01	0.02
Furn. & dur. HH eqpt	0.11	0.02	-0.06 (0.01)	0.08	0.06	0.02	0.02	0.00	0.10	0.00	0.00
Recreat. goods & vehicles	0.20	0.02	-0.06 (0.01)	0.21	0.15	0.06	0.02	0.04	0.24	0.04	0.01
Other durables	0.05	0.01	-0.03 (0.01)	0.05	0.03	0.02	0.00	0.01	0.06	0.01	0.00
Goods, nondurable	0.32	0.15	-0.19 (0.03)	0.38	0.20	0.18	0.13	0.05	0.37	0.05	0.01
Food & beve. off premises	0.06	0.05	-0.03 (0.02)	0.08	0.03	0.04	0.03	0.01	0.08	0.02	0.00
Clothing & footwear	0.06	0.02	-0.05 (0.01)	0.12	0.03	0.09	0.07	0.02	0.08	0.02	0.01
Gasoline & energy	0.00	0.02	-0.03 (0.01)	0.03	-0.02	0.04	0.02	0.02	0.00	0.00	0.01
Other nondurable goods	0.19	0.06	-0.07 (0.01)	0.15	0.15	0.00	0.00	0.00	0.19	0.00	0.01
Services	0.76	0.46	-0.30 (0.08)	1.21	0.57	0.64	0.46	0.18	0.93	0.18	0.02
Housing & utilities	0.13	0.13	-0.06 (0.02)	0.28	0.10	0.18	0.12	0.07	0.20	0.07	0.01
Health care	0.31	0.11	0.00 (0.03)	0.23	0.31	-0.08	-0.07	-0.01	0.31	0.00	0.01
Transportation services	0.04	0.02	-0.08 (0.01)	0.07	0.00	0.07	0.06	0.01	0.03	-0.01	0.00
Recreational services	0.04	0.03	-0.04 (0.01)	0.09	0.02	0.07	0.05	0.03	0.06	0.02	0.00
Food serv. & accomm.	0.11	0.04	-0.06 (0.02)	0.09	0.08	0.01	0.01	0.00	0.11	0.00	0.01
Fin. services & insurance	0.00	0.05	-0.02 (0.04)	0.18	-0.03	0.21	0.16	0.05	0.06	0.07	0.01
Other services	0.10	0.06	-0.06 (0.02)	0.15	0.06	0.09	0.07	0.02	0.11	0.01	0.01
NPISH	0.03	0.02	0.02 (0.01)	0.12	0.05	0.07	0.05	0.02	0.06	0.03	0.01
Gross priv. dom. investment	0.91	0.15	-1.11 (0.14)	0.63	0.03	0.60	0.45	0.15	0.89	-0.02	0.04
Fixed private investment	0.70	0.15	-0.94 (0.07)	0.53	0.09	0.43	0.41	0.03	0.59	-0.11	0.03
Nonresidential	0.50	0.12	-0.69 (0.08)	0.47	0.13	0.34	0.26	0.08	0.48	-0.02	0.02
Structures	-0.01	0.03	-0.19 (0.03)	-0.01	-0.06	0.05	0.02	0.03	0.00	0.01	0.01
Equipment	0.38	0.06	-0.44 (0.05)	0.30	0.09	0.20	0.17	0.03	0.33	-0.05	0.02
Intell. property products	0.14	0.04	-0.06 (0.01)	0.19	0.11	0.07	0.06	0.02	0.15	0.01	0.01
Residential	0.20	0.03	-0.25 (0.05)	0.07	-0.03	0.10	0.15	-0.05	0.12	-0.08	0.02
Structures	0.20	0.03	-0.25 (0.05)	0.07	-0.03	0.10	0.15	-0.05	0.11	-0.08	0.02
Equipment	0.00	0	0.00 (0.00)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Government expenditures	-0.19	0.19	0.10 (0.06)	0.45	-0.11	0.56	0.31	0.25	0.12	0.31	0.03
Federal	-0.09	0.08	0.11 (0.05)	0.18	-0.04	0.22	0.07	0.15	0.11	0.20	0.02
State & local	-0.10	0.12	-0.01 (0.03)	0.26	-0.07	0.34	0.24	0.10	0.02	0.12	0.01
Exports	0.58	0.13	-0.27 (0.08)	0.60	0.36	0.24	0.10	0.14	0.60	0.03	0.04
Imports	-0.76	-0.16	0.54 (0.09)	-0.70	-0.34	-0.36	-0.29	-0.08	-0.77	-0.01	0.03
Addendum:											
Government cons. expend. + transfer payments	0.66		1.22 (0.52)	3.67	1.33	2.34	0.91	1.44	2.86	2.20	0.23

Notes: Indented components add to the final entry at the prior level of indentation.

Table 6: Test Statistics for a Break in the Mean Growth Rate in TFP

	QLR (sup-Wald) test 1 break	2 breaks	3 breaks	Nyblom test	LFST test
A. 1956-2016					
p-value for H_0: $\mu_t = \mu$	0.01	0.06	0.01	0.02	0.03
Estimated break dates	1973Q1	1973Q1 2006Q1	1973Q1 1995Q4 2006Q1		
$\hat{\sigma}_{\Delta\mu}$	0.11			0.11	
90% CI for $\sigma_{\Delta\mu}$	(0.03, 0.36)			(0.02, 0.40)	
B. 1981-2016					
p-value for H_0: $\mu_t = \mu$	0.38	0.14	0.25	0.35	0.31
Estimated break dates	2006Q1	1995Q1 2006Q1	1988Q1 1995Q4 2006Q1		
$\hat{\sigma}_{\Delta\mu}$	0.05			0.05	
90% CI for $\sigma_{\Delta\mu}$	(0.0, 0.15)			(0.0, 0.27)	

Notes: All test are of a constant mean against a non-constant alternative: for the QLR, regime changes; for the Nyblom, against random walk drift; for the LFST, against more general martingale variation. All tests are heteroskedasticity and autocorrelation-robust. The final two rows in each block provide the point estimate of the standard deviation of a random walk drift in the mean, $\sigma_{\Delta\mu}$, and its 90% confidence interval based on inverting the test statistic.

Table 7: Industry Growth by Subperiod

		Pre-1995	1995-2000	2000-2004	2004-2007	2007-2014	Change after 2004 (d-c)	VA Weight
		(a)	(b)	(c)	(d)	(e)	(f)	(g)
(1)	**Private business**	**0.62**	**1.35**	**2.05**	**0.32**	**0.26**	**-1.73**	100.0
(2)	Finance and Insurance	-0.48	3.37	0.89	0.27	0.22	-0.63	8.3
(3)	Energy (Oil/gas, pipeline, refining, utilities)	3.15	-3.47	5.55	-3.51	3.14	-9.06	5.9
(4)	Transportation (ex. pipelines)	3.47	2.34	2.57	2.78	0.40	0.21	2.5
(5)	Construction	0.17	-1.29	-0.82	-5.50	-0.62	-4.67	6.0
(6)	IT producing	8.47	14.46	7.23	6.78	2.49	-0.45	5.7
(7)	Business ex. finance	0.71	1.17	2.17	0.34	0.28	-1.84	91.7
(8)	Finance intensive	0.22	0.24	1.35	-0.03	0.57	-1.37	44.7
(9)	Non-finance intensive	1.16	2.03	2.95	0.67	-0.03	-2.28	47.0
(10)	Business ex. finance and IT prod	0.25	0.23	1.84	-0.10	0.12	-1.93	86.0
(11)	IT-intensive	0.39	0.96	2.19	0.86	-0.22	-1.33	42.8
(12)	Non-IT-intensive	0.11	-0.52	1.49	-0.99	0.45	-2.49	43.2

Notes: Industry and aggregate growth based on BLS 60-industry MFP data. Entries are percent change per year, except for value-added weight, which is average percentage share from 1988-2014.

Table 8: Panel Regressions of Industry TFP Growth on Regulatory Restrictions

	(1)	(2)	(3)	(4)
$Regulation_{i,t}$	0.032 (0.032)	0.033 (0.033)		
$Regulation_{i,t\text{-}1}$	-0.023 (0.027)	-0.011 (0.026)		
$Regulation_{i,t\text{-}2}$	-0.045 (0.039)	-0.036 (0.035)		
$Regulation_{i,t\text{-}3}$	0.022 (0.023)	0.036 (0.034)		
$\overline{Regulation}_{i,t:t-2}$			-0.018 (0.040)	-0.009 (0.036)
$\overline{Regulation}_{i,t-3:t-5}$				0.060 (0.050)
Industry Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	No	Yes	Yes	Yes
F-statistics for *Regulation* (p-value)	0.36 (0.83)	0.44 (0.78)	0.19 (0.67)	0.86 (0.43)

Notes: Data are annual observations of industry TFP growth (the dependent variable) and regulations for the 42 industries for which Regdata has an index of regulation, 1988-2014. Standard errors (in parentheses) are clustered by industry. $\overline{Regulation}_{i,t:t-2}$ denotes the average value of *Regulation* for lags 0-2, and $\overline{Regulation}_{i,t-3:t-5}$ is defined analogously.

Table 9: Changes in Weekly Hours of Time Use, 2007 to 2015, People 15 and Older

	Personal care, including sleep	*Market work*	*Education*	*Leisure*	*Other*
Men	2.0	-2.4	0.5	0.7	-0.8
Women	2.4	-1.5	0.1	0.7	0.8

Fernald, Hall, Stock, and Watson
March 10, 2017

Supplementary Econometric Appendix

Estimation of the trend and its standard error. As discussed in Section II, we estimate μ_t using the partially linear regression model. An alternative would be to make an explicit parametric assumption about the process followed by μ_t and z_t. For example, Gordon (2014) estimates cyclically adjusted trends by modeling μ_t as a Gaussian random walk and z_t as serially uncorrelated and Gaussian, with μ and z being independent. This gives a fully specified likelihood that can be maximized using state space methods. Gordon (2014) then estimates μ_t using the Kalman smoother. Although these two approaches sound different, in the end they both involve estimation of μ_t by smoothing $y_t - \hat{\beta}(L)\Delta u_t$ as in (9). Stock and Watson (2016, Figure 2) compare the lag weights for the biweight filter and the implied filter from the Kalman smoother for the random-walk model of μ_t. On a series-by-series basis, the Kalman smoother and partially linear regression approaches often give quite similar results. However, because the state-space approach entails estimation of different model parameters for each series, the implied smoothing filters differ across series so the additivity property discussed in the next subsection does not hold for the state space approach.

The use of regression (9) to estimate the trend departs from standard practice in partially linear regression, in which β(L) is estimated by regressing prefiltered (1-κ(L))y_t on leads and lags of prefiltered (1-κ(L))Δu_t, but this departure is justified theoretically when the variation in μ_t is small compared with the variation in Δu_t and z_t, as it is here, and in any event the two estimation methods yield virtually identical results. We use the simple approach here for transparency and to stay as close as possible to conventional implementations of Okun's Law.

HAC standard errors for $\hat{\mu}_t$ are computed as follows. With some abuse of notation, write $\hat{\mu}_t = \kappa_t'\nu$, where κ_t is a T-vector of weights associated with κ(L) and ν is the T-vector with $\nu_t = y_t - \hat{\beta}(L)\Delta u_t$. Then $\text{var}(\hat{\mu}_t) = \text{var}(\kappa_t'\nu) = \kappa_t'\Sigma_\nu\kappa_t$, where Σ_ν is the $T\times T$ covariance matrix of ν. If κ_t were the vector of ones, then $\text{var}(\hat{\mu}_t)$ is the HAC estimation problem of estimating the variance of the mean. The problem here is closely related in that κ_t has many very similar values. There are many ways to address the HAC problem. Here, we chose a simple method for reliably computing positive semidefinite inner products by approximating the stochastic process for ν_t as a first order autoregression, then estimating $\kappa_t'\Sigma_\nu\kappa_t$ using the implied parametric covariance matrix.

Computation of the factor forecasts. The factor forecasts described in Section II.B are computed as follows.

(i) All 123 series used to estimate the factors are transformed to approximate stationarity. Real activity variables are transformed to (annualized) growth rates, inflation is transformed to first differences, interest rates and unemployment rates appear in first differences, spreads and ratios that are approximately cointegrating appear as differences of levels or log levels. (The specific transformation applied to each series is listed in the data appendix.) Any remaining near-zero frequency variation is removed by local demeaning using a biweight kernel with 25-year bandwidth. See Stock and Watson (2016).

(ii) The factors are estimated by principal components (computing using least squares on the unbalanced panel of data) over the period 1959, third quarter through 2016, second quarter.

(iii) The DFM parameters Λ and Φ(L), with 4 factors and 4 VAR lags, are estimated by OLS, using data from 1984, first quarter, to 2009, second quarter, treating the estimated factors $\hat{F}_t$ as data. The start date of 1984, first quarter, is chosen to align with standard estimates of the start of the Great Moderation period. There is evidence of a break in the factor loadings around this date, see Stock and Watson (2016) for a review of this literature. As discussed there, even if there are structural breaks in the dynamic factor model coefficients it can be desirable to estimate the factors over the full sample (here 1959-2016), and this appears to be the case for this data set.

(iv) Given the factors through the trough quarter, forecasts of the factors, $\hat{F}_{t|2009q2}$, are computed for succeeding quarters using the factor vector autoregression and history of the factors through the 2009 trough date.

(v) Given the factor forecasts, forecasts of the detrended variables are computed as $X_{t|2009q2} = \hat{\Lambda}\hat{F}_{t|2009q2}$ for the succeeding period, where the estimated value of Λ is computed by using data from 1984, second quarter, through the 2009 trough date.

(vi) Forecasts for the original series (not detrended) are computed by adding the forecast of the detrended variable to the trough value of the trend, adjusted as appropriate for the demographic trend in the labor force participation rate (details discussed below). The results are robust to variations in the benchmark model, including shifting the jumping-off date to 2009, fourth quarter. For the main series, including output and employment, and productivity, they are also robust to using a low-dimensional vector autoregression.

In the context of the trend-cycle-irregular decomposition, the forecast from (v) is the estimated cyclical component of the series c_t, the forecast from (vi) is the estimated trend + cyclical component $\mu_t + c_t$, and the forecast error—the unexpected shortfall or exceedance of y_t—is the irregular component z_t.

Section V estimated the decline in the labor force participation rate (LFPR) associated with changing demographics in the population. Because these changing demographics were largely known or could have been accurately forecast over the 2009-2016 period, we incorporate these changes in the forecast of the growth rate for LFPR. These demographic adjustments to the LFPR growth rate forecast are included one-for-one in the forecast growth rate of employment and hours, and in the various output measures after multiplying by labor's share. Forecasts for the trends in capital, TFP, and labor quality are left unchanged. To maintain adding-up for the expenditure decomposition of GDP in Table 5, the trend in each expenditure component is adjusted by its share in GDP multiplied by the LFPR demographic trend adjustment in GDP.

Bayesian Implementation of random-walk-plus noise model for TFP growth rates. Figure 11 and Figure 12 show results from a model for TVP growth in which the growth rate of TFP, say y_t, follows the model $y_t = \beta(L)\Delta u_t + \mu_t + z_t$, where $\Delta\mu_t = \eta_t$ and $\{z_t\}$ and $\{\eta_t\}$ are mutually independent Gaussian white noise processes that are independent of Δu_t. We fixed $\beta(L)$ at its OLS estimate and estimated σ_z, σ_η, and the time path of $\{\mu_t\}$ using Bayes methods using independent priors for σ_z, σ_η, and μ_0. Specifically, $\mu_0 \sim N(1,10)$, $\sigma_z \sim u[0.67s,1.33s]$ where s is the sample standard deviation of $y_t - \beta(L)\Delta u_t$, and $\sigma_\eta \sim U[0,0.25]$. Posteriors for $\{\mu_t\}$, σ_z and σ_η, were computed using $y_t - \beta(L)\Delta u_t$ for $t \in [1956{:}q3,2016{:}q2]$ and $t \in [1981{:}q3,2016{:}q2]$. Posterior quantiles for $\{\mu_t\}$ are shown in Figure 11 based on the 1956-2016 sample. The marginal posteriors for σ_z and $\sigma_{\Delta\eta}$ are shown in Appendix Figure 1.

Appendix Figure 1:

Priors and Posteriors for Random-walk + white noise model for TFP growth rates

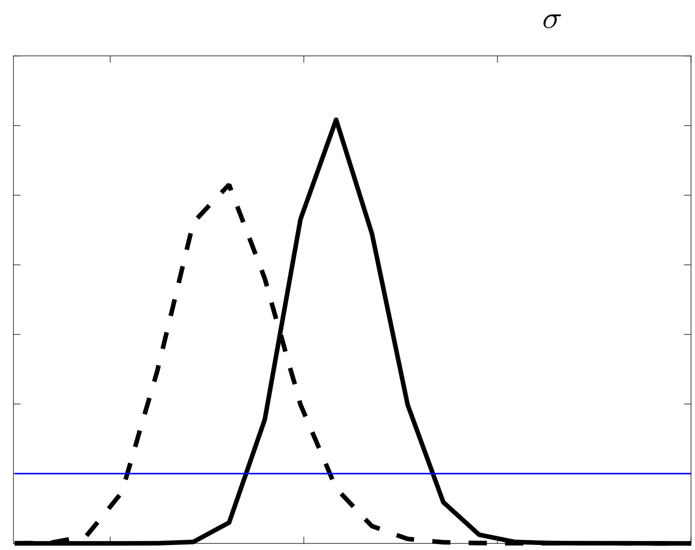

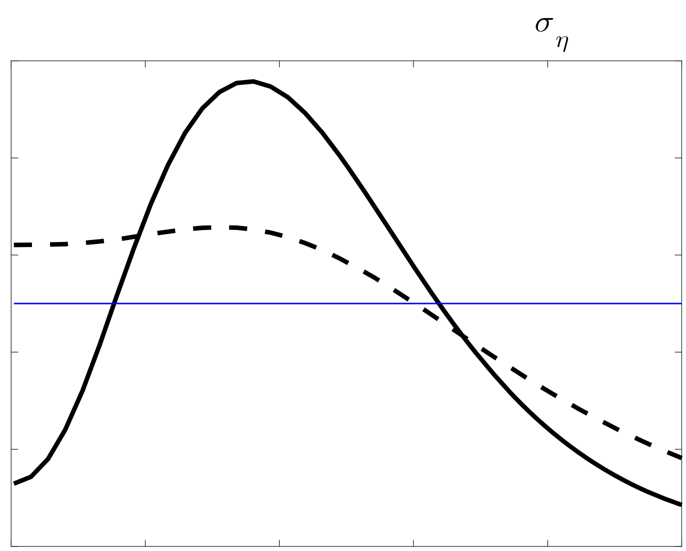

Notes: 1956-2016 posterior (solid black), 1981-2016 posterior (dashed), prior (thin solid blue).

Variables used to construct factors for the DFM

The DFM factors were estimated using principal methods surveyed in Stock and Watson (2015) and 123 time series from an updated version of the dataset described in that paper. The variables are from 12 broad categories shown in table 1. The specific series are listed below.

Table A.1: Data Series

	Name	Description	Sample Period	T
(1) NIPA				
1	Cons:Dur	Real personal consumption expenditures: Durable goods	1959:Q1, 2016:Q3	5
2	Cons:Svc	Real personal consumption expenditures: Services	1959:Q1, 2016:Q3	5
3	Cons:NonDur	Real personal consumption expenditures: Nondurable goods	1959:Q1, 2016:Q3	5
4	Inv:Equip	Real Gross Private Domestic Investment: Fixed Investment: Nonresidential: Equipment	1959:Q1, 2016:Q3	5
5	FixInv:NonRes	Real private fixed investment: Nonresidential	1959:Q1, 2016:Q3	5
6	FixedInv:Res	Real private fixed investment: Residential	1959:Q1, 2016:Q3	5
7	Ch. Inv/GDP	Change in Inventories /GDP	1959:Q1, 2016:Q3	1
8	Gov:Fed	Real government consumption expenditures and gross investment: Federal	1959:Q1, 2016:Q3	5
9	Real_Gov Receipts	Government Current Receipts (Nominal) Defl by GDP Def	1959:Q1, 2016:Q3	5
10	Gov:State&Local	Real government consumption expenditures and gross investment: State and local	1959:Q1, 2016:Q3	5
11	Exports	Real exports of goods and services	1959:Q1, 2016:Q3	5
12	Imports	Real imports of goods and services	1959:Q1, 2016:Q3	5
(2) Industrial Production				
13	IP: Dur gds materials	Industrial Production: Durable Materials	1959:Q1, 2016:Q3	5
14	IP: Nondur gds materials	Industrial Production: nondurable Materials	1959:Q1, 2016:Q3	5
15	IP: Dur Cons. Goods	Industrial Production: Durable Consumer Goods	1959:Q1, 2016:Q3	5
16	IP: Auto	IP: Automotive products	1959:Q1, 2016:Q3	5
17	IP:NonDur Cons God	Industrial Production: Nondurable Consumer Goods	1959:Q1, 2016:Q3	5
18	IP: Equip	Industrial Production: Equipment, total, Index 2012=100, Monthly, Seasonally Adjusted	1959:Q1, 2016:Q3	5
19	Capu Tot	Capacity Utilization: Total Industry	1967:Q1, 2016:Q3	1
(3) Employment and Unemployment				
20	Emp: DurGoods	All Employees: Durable Goods Manufacturing	1959:Q1, 2016:Q3	5
21	Emp: Const	All Employees: Construction	1959:Q1, 2016:Q3	5
22	Emp: Edu&Health	All Employees: Education & Health Services	1959:Q1, 2016:Q3	5
23	Emp: Finance	All Employees: Financial Activities	1959:Q1, 2016:Q3	5
24	Emp: Infor	All Employees: Information Services	1959:Q1, 2016:Q3	5
25	Emp: Bus Serv	All Employees: Professional & Business Services	1959:Q1, 2016:Q3	5
26	Emp:Leisure	All Employees: Leisure & Hospitality	1959:Q1, 2016:Q3	5
27	Emp:OtherSvcs	All Employees: Other Services	1959:Q1, 2016:Q3	5
28	Emp: Mining/NatRes	All Employees: Natural Resources & Mining	1959:Q1, 2016:Q3	5
29	Emp:Trade&Trans	All Employees: Trade Transportation & Utilities	1959:Q1, 2016:Q3	5
30	Emp:Retail	All Employees: Retail Trade	1959:Q1, 2016:Q3	5
31	Emp:Wholesal	All Employees: Wholesale Trade	1959:Q1, 2016:Q3	5
32	Emp: Gov(Fed)	Employment Federal Government	1959:Q1, 2016:Q3	5
33	Emp: Gov (State)	Employment State government	1959:Q1, 2016:Q3	5
34	Emp: Gov (Local)	Employment Local government	1959:Q1, 2016:Q3	5
35	Urate: Age16-19	Unemployment Rate - 16-19 yrs	1959:Q1, 2016:Q3	2

36	Urate:Age>20 Men	Unemployment Rate - 20 yrs. & over Men	1959:Q1, 2016:Q3	2
37	Urate: Age>20 Women	Unemployment Rate - 20 yrs. & over Women	1959:Q1, 2016:Q3	2
38	U: Dur<5wks	Number Unemployed for Less than 5 Weeks	1959:Q1, 2016:Q3	5
39	U:Dur5-14wks	Number Unemployed for 5-14 Weeks	1959:Q1, 2016:Q3	5
40	U:dur>15-26wks	Civilians Unemployed for 15-26 Weeks	1959:Q1, 2016:Q3	5
41	U: Dur>27wks	Number Unemployed for 27 Weeks & over	1959:Q1, 2016:Q3	5
42	U: Job losers	Unemployment Level - Job Losers	1967:Q1, 2016:Q3	5
43	U: LF Reenty	Unemployment Level - Reentrants to Labor Force	1967:Q1, 2016:Q3	5
44	U: Job Leavers	Unemployment Level - Job Leavers	1967:Q1, 2016:Q3	5
45	U: New Entrants	Unemployment Level - New Entrants	1967:Q1, 2016:Q3	5
46	Emp:SlackWk	Employment Level - Part-Time for Economic Reasons All Industries	1959:Q1, 2016:Q3	5
47	AWH Man	Average Weekly Hours: Manufacturing	1959:Q1, 2016:Q3	1
48	AWH Privat	Average Weekly Hours: Total Private Industrie	1964:Q1, 2016:Q3	2
49	AWH Overtime	Average Weekly Hours: Overtime: Manufacturing	1959:Q1, 2016:Q3	2
		(4) Orders, Inventories and Sales		
50	Orders:Dur Goods	New Orders for Durable Goods Defl by PCE(LFE) Def	1959:Q1, 2016:Q3	5
51	Orders:ConsGoods	New Orders for Consumer Goods Defl by PCE(LFE) Def	1992:Q1, 2016:Q2	5
52	Unfilledorders	Unfilled Orders for Durable Goods Defl by PCE(LFE) Def	1959:Q1, 2016:Q3	5
53	Orders:capgds	New Orders for Nondefense Capital Goods Defl by PCE(LFE) Def	1968:Q1, 2016:Q3	5
54	VendPerf	ISM Manufacturing: Supplier Deliveries Index©	1959:Q1, 2014:Q4	1
55	NAPM:ORD	ISM Manufacturing: New Orders Index©; Index;	1959:Q1, 2014:Q4	1
56	Business Inventory	Total Business Inventories Defl by PCE(LFE) Def	1959:Q1, 2016:Q2	5
57	Inv/Sales	Total Business: Inventories to Sales Ratio	1959:Q1, 2016:Q2	2
		(5) Housing Starts and Permits		
58	Hpermits	New Private Housing Units Authorized by Building Permit	1960:Q1, 2016:Q3	5
59	Hstarts:MW	Housing Starts in Midwest Census Region	1959:Q1, 2016:Q3	5
60	Hstarts:NE	Housing Starts in Northeast Census Region	1959:Q1, 2016:Q3	5
61	Hstarts:S	Housing Starts in South Census Region	1959:Q1, 2016:Q3	5
62	Hstarts:W	Housing Starts in West Census Region	1959:Q1, 2016:Q3	5
63	Constr. Contracts	Construction contracts (mil. sq. ft.) (Copyright McGraw-Hill)	1963:Q1, 2014:Q4	4
		(6) Prices		
64	GPDI Defl	Gross Private Domestic Investment: Chain-type Price Index	1959:Q1, 2016:Q3	6
65	BusSec Defl	Business Sector: Implicit Price Deflator	1959:Q1, 2016:Q3	6
66	PCED_MotorVec	Motor vehicles and parts	1959:Q1, 2016:Q3	6
67	PCED_DurHousehold	Furnishings and durable household equipment	1959:Q1, 2016:Q3	6
68	PCED_Recreation	Recreational goods and vehicles	1959:Q1, 2016:Q3	6
69	PCED_OthDurGds	Other durable goods	1959:Q1, 2016:Q3	6
70	PCED_Food_Bev	Food and beverages purchased for off-premises consumption	1959:Q1, 2016:Q3	6
71	PCED_Clothing	Clothing and footwear	1959:Q1, 2016:Q3	6
72	PCED_Gas_Enrgy	Gasoline and other energy goods	1959:Q1, 2016:Q3	6
73	PCED_OthNDurGds	Other nondurable goods	1959:Q1, 2016:Q3	6
74	PCED_Housing-Utilities	Housing and utilities	1959:Q1, 2016:Q3	6
75	PCED_HealthCare	Health care	1959:Q1, 2016:Q3	6
76	PCED_TransSvg	Transportation services	1959:Q1, 2016:Q3	6
77	PCED_RecServices	Recreation services	1959:Q1, 2016:Q3	6
78	PCED_FoodServ_Acc.	Food services and accommodations	1959:Q1, 2016:Q3	6
79	PCED_FIRE	Financial services and insurance	1959:Q1, 2016:Q3	6
80	PCED_OtherServices	Other services	1959:Q1, 2016:Q3	6
81	PPI:FinConsGds	Producer Price Index: Finished Consumer Goods	1959:Q1, 2015:Q4	6

82	PPI:FinConsGds(Food)	Producer Price Index: Finished Consumer Foods	1959:Q1, 2015:Q4	6
83	PPI:IndCom	Producer Price Index: Industrial Commodities	1959:Q1, 2016:Q3	6
84	PPI:IntMat	Producer Price Index: Intermediate Materials: Supplies & Components	1959:Q1, 2015:Q4	6
85	P:SensMat	Index of Sensitive Matrerials Prices (Discontinued) Defl by PCE(LFE) Def	1959:Q1, 2004:Q1	5
86	NAPM com price	ISM Manufacturing: Prices Paid Index©	1959:Q1, 2014:Q4	1
87	Price:NatGas	PPI: Natural Gas Defl by PCE(LFE) Def	1967:Q1, 2016:Q3	5
		(7) Productivity and Earnings		
88	CPH:NFB	Nonfarm Business Sector: Real Compensation Per Hour	1959:Q1, 2016:Q3	5
89	CPH:Bus	Business Sector: Real Compensation Per Hour	1959:Q1, 2016:Q3	5
90	OPH:nfb	Nonfarm Business Sector: Output Per Hour of All Persons	1959:Q1, 2016:Q3	5
91	ULC:NFB	Nonfarm Business Sector: Unit Labor Cost	1959:Q1, 2016:Q3	5
92	UNLPay:nfb	Nonfarm Business Sector: Unit Nonlabor Payments	1959:Q1, 2016:Q3	5
		(8) Interest Rates		
93	FedFunds	Effective Federal Funds Rate	1959:Q1, 2016:Q3	2
94	TB-3Mth	3-Month Treasury Bill: Secondary Market Rate	1959:Q1, 2016:Q3	2
95	BAA_GS10	BAA-GS10 Spread	1959:Q1, 2016:Q3	1
96	MRTG_GS10	Mortg-GS10 Spread	1971:Q2-2016:Q3	1
97	tb6m_tb3m	tb6m-tb3m	1959:Q1, 2016:Q3	1
98	GS1_tb3m	GS1_Tb3m	1959:Q1, 2016:Q3	1
99	GS10_tb3m	GS10_Tb3m	1959:Q1, 2016:Q3	1
100	CP_Tbill Spread	CP3FM-TB3MS	1959:Q1, 2016:Q3	1
101	Ted_spr	MED3-TB3MS (Version of TED Spread)	1971:Q1, 2016:Q3	1
		(9) Credit		
102	C&L loans	Commercial and Industrial Loans at All Commercial Banks Defl by PCE(LFE) Def	1959:Q1, 2016:Q3	5
103	ConsLoans	Consumer (Individual) Loans at All Commercial Banks, adjusted for outlier in April 2010 (see FRB H8 Release) Defl by PCE(LFE) Def	1959:Q1, 2016:Q3	5
104	NonRevCredit	Total Nonrevolving Credit Outstanding Defl by PCE(LFE) Def	1959:Q1, 2016:Q3	5
105	LoansRealEst	Real Estate Loans at All Commercial Banks Defl by PCE(LFE) Def	1959:Q1, 2016:Q3	5
106	RevolvCredit	Total Revolving Credit Outstanding Defl by PCE(LFE) Def	1968:Q1, 2016:Q3	5
		(10) Exchange Rates		
107	Ex rate: major	FRB Nominal Major Currencies Dollar Index (Linked to EXRUS in 1973:1)	1959:Q1, 2016:Q3	5
108	Ex rate: Euro	U.S. / Euro Foreign Exchange Rate	1999:Q1, 2016:Q3	5
		(11) Asset Prices, Wealth, and Household Balance Sheets		
109	S&P 500	S&P's Common Stock Price Index: Composite (1941, 43=10)	1959:Q1, 2016:Q3	5
110	HHW:TL	Real Total Liabilities of Households and Non Profits (billions of $2009) deflated by core PCE .. Fred-QD. Seasonally adjusted using RATS-X11	1959:Q1, 2016:Q3	5
111	HHW:W	Real Net Worth of Households and Non profits (billions of $2009) deflated by core PCE .. FREDQD. Seasonally adjusted using RATS-X11	1959:Q1, 2016:Q3	5
112	HHW:TA_XRE	Real Assets of households and nonprofits, excluding real estate (billions of $2009) def. by core PCE, FredQD . Seasonally adjusted using RATS-X11	1959:Q1, 2016:Q3	5
113	HHW:TA_RE	Real Real Estate Assets of households and and Nonprofits (billions of $2009) defl by core PCE .. FREDQD. Seasonally adjusted using RATS-X11	1959:Q1, 2016:Q3	5

114	DJIA	Common stock prices: Dow Jones industrial average	1959:Q1, 2016:Q3	5
115	VXO	VXO	1962:Q3-2016:Q3	1
116	CS_10	Case-Shiller 10 City Average Defl by PCE(LFE) Def	1987:Q1, 2016:Q3	5
117	CS_20	Case-Shiller 20 City Average Defl by PCE(LFE) Def	2000:Q1, 2016:Q3	5
(12) Asset Prices				
118	IP: Energy Prds	IP: Consumer Energy Products	1959:Q1, 2016:Q3	5
119	Price:Oil	PPI: Crude Petroleum Defl by PCE(LFE) Def	1959:Q1, 2016:Q3	5
120	Crudeoil Price	Crude Oil: West Texas Intermediate (WTI) - Cushing Oklahoma Defl by PCE(LFE) Def	1986:Q1, 2016:Q3	5
121	CrudeOil	Crude Oil Prices: Brent - Europe Defl by PCE(LFE) Def	1987:Q3-2016:Q3	5
122	Price Gasoline	Conventional Gasoline Prices: New York Harbor Regular Defl by PCE(LFE) Def	1986:Q3-2016:Q3	5
123	CPI Gasoline	CPI Gasoline (NSA) BLS: CUUR0000SETB01 Defl by PCE(LFE) Def	1959:Q1, 2016:Q3	5

Notes: The final column “T” indicates how the variable was transformed 1 = no transformation; 2 = first difference; 3 = second difference; 4 = logarithm; 5 = first difference of logarithm; 6 = second difference of logarithm.