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 Weak Recovery of Output

February 6, 2017 draft

ABSTRACT: The U.S. economy has been expanding since the recession trough in 2009. Though unemployment has declined at about the same rate as in previous recoveries, output has grown much more slowly than in the past. We explore explanations for the shortfall in output growth, using a quantitative decomposition based on growth economics. Two components of the decomposition stand out: slow growth in productivity, and a growing shortfall of labor-force participation relative to its demographic determinants. The slow growth in both components predated the recession. Our analysis gives a full treatment to cyclical effects. Our conclusion is that powerful non-cyclical forces at least partially unrelated to the financial crisis of 2008 account for the poor record of output growth during the ongoing recovery from the crisis-induced recession. This combination of the adverse cyclical influence, and the noncyclical forces we study, resulted in a shortfall of capital formation that holds back output even today.

We thank Lucretia Reichlin, Larry Ball, and other participants at the October 2016 Federal Reserve Bank of Boston conference on “The Elusive ‘Great’ Recovery.” We also thank Neil Gerstein for excellent research assistance. The views expressed in this paper are those of the authors and should not be attributed to the Federal Reserve Bank of San Francisco, the Federal Reserve System’s Board of Governors, or the Federal Reserve System, or the Business Cycle Dating Committee of the National Bureau of Economic Research.

Output growth since the trough of the 2007-2009 recession has fallen far short of both forecasts and past experiences in cyclical expansions. Output per capita grew at only about a one percent annual pace from mid-2009 through mid-2016, by far the slowest expansionary pace in the past 70 years and more than a percentage point below its average during the three preceding expansions. During this anemic episode of output growth, employment has grown at closer to historical rates and, a s can be seen in Figure 1, the unemployment rate has fallen at least as rapidly as in recent expansions. The unemployment rate, which peaked at 10 percent in late 2009, fell below 5 percent by early 2016.

This slow growth in output was unexpected. Figure 2 shows the median forecast paths of GDP and the unemployment rate, from the Survey of Professional Forecasters, for forecasts made annually in the first-quarter 2010 through 2015 using data through the end of the previous year. These median forecasts were consistently overly optimistic about GDP growth. Early in the recovery, forecasts of the decline in the unemployment rate were borne out, but starting in 2013 they were too pessimistic. These forecasts, which are representative of other real-time forecasts by the Congressional Budget Office, the Federal Reserve Board, and the Council of Economic Advisers, pose the core challenge tackled in this paper: Why was output growth so slow in this recovery, given the normal, or better-than-normal, recovery in the labor market?

To address this question, we begin by using Robert Solow's growth-accounting decomposition to quantify the slowdown and to decompose the slow growth into its sources. Output per capita grew about 1.3 percent more slowly in the expansion starting in 2009 than in the three previous expansions, those following the recessions of 1981-82, 1990-91, and 2001. This direct comparison, however, understates the shortfall in growth because of the severity of the 2007-2009 recession and the large subsequent decline in the unemployment rate during the expansion—which, in past experience, would be associated with rapid growth. We therefore use two methods to adjust for the severity of the recession. The first method conditions on the unemployment path during the recovery, and uses Okun's law to impute a normal cyclical recovery; this provides a counterfactual basis which allows estimating the output growth shortfall, given the strength of the recovery in the labor market. The second method conditions on the state of the economy at the cyclical trough of 2009q2; by doing so it asks the different, but related, question of what a normal path of output and employment would be, given the severity of the recession as observed at the trough. This is effectively a forecasting question, so the second method examines pseudo out-of-sample forecasts using a many-variable time series model.

Given the strength of the recovery in the labor market, the growth in output is even more disappointing: conditional on the path of the unemployment rate, the cyclically adjusted growth rate of output per capita was 1.8 percent per year slower than the average of the adjusted growth in the three previous expansions. Cumulatively, this amounts to an enormous 14 percent shortfall in the level of output per capita over the first seven years of the recovery.

Our growth accounting attributes the shortfall in cyclically adjusted growth primarily to declining rates of labor force participation and to slow growth in (total factor) productivity, both cyclically adjusted. In a proximate sense, lack of capital deepening is also important. But we tend to interpret the decline in participation and in productivity growth as exogenous and capital formation as endogenous. Observed capital growth is consistent with the observed pace of growth in productivity and labor input, in that the capital-output ratio has largely tracked its historical experience. Thus, we do not put the shortfall in capital formation in our list of more exogenous causal factors, but interpret it mainly as a response to the two major adverse forces. Other growth-accounting factors are quantitatively minor.

As we discuss, both declining participation and slow growth in productivity predate the recession itself, and appear to be largely independent of other economic developments, including the financial crisis. After adjusting for cyclical variation, the decline in participation largely follows patterns that have evolved at long time scales, notably the surge of women into the labor force in the 1970s and 1980s, the aging of the baby boom generation, and the persistent long-term decline in prime-age male participation (which since 2000 has been mirrored in prime-age female participation). Although much of the literature on the recent decline in the labor force participation rate focuses on the retirement of the baby boom, we provide new evidence that suggests the numerical contribution of this effect is less clear cut, and to date perhaps smaller, than sometimes suspected. In any event, there is no dispute that the decline in participation, whatever its source, has been a drag on output, and largely reflects long-term trends that predate the 2007-2009 recession.

The other main source of the slow growth, in a growth accounting sense, is the disappointing growth in productivity. Throughout the paper, “productivity” means total factor productivity, the natural measure of productivity in Solow’s framework. We will be specific about our use of other concepts of productivity, such as labor productivity (output per hour of labor input). Relative to its trend from the early 1970s to the mid-1990s, productivity growth rose markedly in the mid-1990s, but slowed after the mid-2000s. Since 2004 or so, the pace of productivity growth has been similar to its pace in the 1970s and 1980s. As we discuss more fully below, one view that is consistent with the data is that there was a transformative wave of productivity growth linked to the invention and use of information technology, but that wave ran its course after the early 2000s. Others believe that the slowdown is only a pause pending the widespread application of artificial intelligence.

Although we focus primarily on understanding these growth-accounting determinants of the slow growth in output – elsewhere referred to as supply-side determinants – there are many explanations of the slow growth that have focused instead on factors that have retarded aggregate demand. Such explanations include consumption growth being retarded because of increasing income inequality, the inability to reduce real interest rates because of the zero lower bound, fiscal drag arising from the sequester in 2013 and 2014, spillovers through weak export demand arising from slow growth in Europe and in China, and the possibility that at least some of the output shortfall is simply a measurement problem and that true output growth is greater than estimated by government statistics. As we discuss below, we find only limited support for many of these explanations, although our estimates do suggest that unusually slow growth in government expenditures – at the federal level, concurrent with the sequester, and at the state and local level, concurrent with the decline in property tax receipts arising from the collapse in house prices during the recession – contributed to the slow growth in demand.

I Growth Decomposition and Data

This section lays out the framework of our application of Solow's growth-accounting framework, and describes the data we use. Section I.A begins with a brief overview of the data and focus. Section I.B discusses the overall Solow-style growth decomposition and Section I.C discusses the variants we use for labor productivity. We focus on expansions as defined by the National Bureau of Economic Research, and compare the values of the variables, averaged over the post-crisis years, to the corresponding average values across the three previous recoveries.

I.A Preliminaries and data

The focus of this paper is the weak recovery in output over the course of the economic recovery since the 2007-2009 recession ended in the second quarter of 2009. Most of our analysis ends in 2016:Q2 because, as explained in Section II, this is the last quarter for which we have the full data for our cyclical decomposition. In some of our analysis, we compare this recovery to the comparable 28-quarter period in the three previous recoveries. For the 2001-2007 recovery, we truncate at the business-cycle peak at the end of 2007 (24 quarters).

The reasons for the weak the recovery can be examined through the lens of production (output must be produced by someone) or expenditure (output is sold to someone). Until Section VI, we focus primarily on the production side, doing a Solow-style growth decomposition of the sources of growth. That is a natural framework for addressing the role of structural trends—

essentially, potential growth—in the slow recovery. We do our main growth accounting at the level of the business economy. The reason is that growth accounting is generally not applicable to government, household, and non-profit activities. For example, much of non-business output is not measured independently of inputs.

Historically, users of the U.S. national accounts generally relied on the expenditure-side measure of output. But recent literature (discussed in the supplemental data appendix) argues for combining it with income-side measures. The idea behind this combination is that 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. Accordingly, throughout the paper, our measure of output is the geometric average of inflation-adjusted income- and product-side measures. At an economy-wide level, we refer to this average of gross domestic product and gross domestic income as gross domestic output. For the rest of the paper, we refer to this variable as output or GDO. Unless noted otherwise, we typically scale output by population aged 16 and above, the population eligible for employment.

We use growth-accounting data for the U.S. business sector described in Fernald (2014). These quarterly data, which are updated regularly and posted on the San Francisco Fed's website, provides the values of the variables in the equations below. Broader real gross product and gross income aggregates were obtained from the Bureau of Economic Analysis, and labor-market data were obtained from the Bureau of Labor Statistics. The supplemental data appendix provides further information.

I.B Accounting for Growth

Although our growth accounting focuses on the business sector, we need to link that to the overall economy. Labor market indicators such as the unemployment rate refer to that concept. Economy-wide output GDO and business output, Y_t^{Bus}, are linked by identities:

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

Pop is the working-age population.

Growth accounting decomposes business-sector output per person into a set of components that we have chosen to help understanding of the evolution of the business economy.[1] Output growth (in log changes) depends on growth in total factor productivity, *productivity*, capital, *K*, and labor, *Labor*. Labor input, in turn, depends on *Hours* and labor quality, *LQ*. In per capita terms:

$$\begin{aligned}\Delta\ln\left(\frac{Y_t^{Bus}}{Pop_t}\right)&=\Delta\ln TFP_t^{Bus}+\alpha_t\Delta\ln\left(\frac{K_t}{Pop_t}\right)+(1-\alpha_t)\Delta\ln\left(\frac{Labor_t^{Bus}}{Pop_t}\right)\\&=\Delta\ln TFP_t^{Bus}+\alpha_t\Delta\ln\left(\frac{K_t}{Pop_t}\right)+(1-\alpha_t)\left(\Delta\ln LQ_t+\Delta\ln\left(\frac{Hours_t^{Bus}}{Pop_t}\right)\right)\end{aligned} \tag{2}$$

The time series α_t is capital's share of income. Under standard conditions, of course, α_t is the elasticity of output with respect to capital. 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 et al. (2016) for a discussion.

Although the growth accounting applies to the business sector, we need to relate business-sector hours to broader labor market indicators. We use the following identity:

$$\left(\frac{Hours_t^{Bus}}{Pop_t}\right)=\left(\frac{Hours_t^{Bus}}{Emp_t^{Bus}}\right)\times\left(\frac{Emp_t^{Bus}}{Emp_t^{CPS}}\right)\times\left(\frac{Emp_t^{CPS}}{LabForce_t}\right)\times\left(\frac{LabForce_t}{Pop_t}\right) \tag{3}$$

The first term on the right-hand side is business-sector hours per employee. The second is the ratio of business employment, measured (primarily) from the establishment survey, to household employment, measured from the Current Population Survey. There are several conceptual differences between the two measures, 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

[1] Solow (1957). Modern growth accounting often follows Jorgenson and Griliches (1967).

business sector, and, correspondingly, is less cyclical.[2] The third term, employment relative to the labor force, is by definition equal to 1-U_t, where U_t is the unemployment rate. The final term is the labor force participation rate, *PARTICIPATION.*

Combining equations (1) and (3) yields a total economy growth-accounting decomposition:

$$\Delta \ln\left(GDO_t / Pop_t\right) = \Delta \ln\left(GDO_t / Y_t^{Bus}\right) + \Delta \ln TFP_t^{Bus} + \alpha_t \Delta \ln\left(K_t / Pop_t\right)$$
$$+ (1-\alpha_t)\left(\Delta \ln LQ_t + \Delta \ln\left(\frac{Hours_t^{Bus}}{Emp_t^{Bus}}\right) + \Delta \ln\left(\frac{Emp_t^{Bus}}{Emp_t^{CPS}}\right) + \Delta \ln\left(1-U_t\right) + \Delta \ln LFPR_t \right) \quad (4)$$

Note that over the long run the contribution of the U term is zero because the unemployment rate reverts fairly predictably to a mean value between 5 and 6 percent.

I.C Labor-productivity decomposition

It is also useful to consider two standard decompositions of business-sector growth in terms of labor productivity. They involve different rearrangements of equation (2):

$$\begin{aligned} \Delta \ln\left(\frac{Y_t^{Bus}}{Hours_t^{Bus}}\right) &= \Delta \ln TFP + \alpha \cdot \Delta \ln\left(\frac{K_t}{LQ_t Hours_t^{Bus}}\right) + \Delta \ln LQ_t \\ &= \frac{\Delta \ln TFP}{(1-\alpha)} + \left(\frac{\alpha}{1-\alpha}\right) \cdot \Delta \ln\left(\frac{K_t}{Y_t}\right) + \Delta \ln LQ_t \end{aligned} \quad (5)$$

These expressions are useful because, in standard growth models, capital deepening is endogenous. For example, in a neoclassical growth model, capital per unit of labor (used in the first line) depends in the long run on technology alone. The decomposition in the second line is useful because, in the one-sector neoclassical model, capital depends on technology and labor in such a way that the capital output ratio is stationary in levels. In more complicated models (and in the data), that ratio has a trend (see, for example, the online appendix to Fernald, 2015). But the capital-output ratio nevertheless gives a natural diagnostic measure of whether weak growth

[2] Fernald and Wang (2016) discuss differences between the business-sector and CPS measures and why this gap is procyclical. (so it tends to fall when the unemployment rate rises).

in capital reflects weak growth in technology and labor (so that both capital and output look weak) or whether there is a large omitted factor.

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, as measured by a large number of time series variables, as of the 2009 trough. 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? Producing the associated counterfactual path amounts to computing a pseudo out-of-sample forecast of multiple time series variables, 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, $y_t = 400\Delta\ln GDO_t$). 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. As explained below, we estimate the long-term trend as the long-run average of y (after subtracting off the cyclical part). This long-run average typically changes over time—for example because of changing demographics. Our quantification finds that those changes are important for understanding the weak recovery.

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 might 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,

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

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, β(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, so we refer to equation (7) as a generalization of Okun's Law, which originally related Δu and to Δy.

The standard unemployment rate, shown in Figure 1, is only one of many measures of the state of the labor force. 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 using these additional measures in an extended version of equation (7), and doing so would be conceptually consistent with our approach. 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 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. We do not follow that approach for two reasons. First, and most importantly, conditioning on output rules out addressing the question of this paper, why output growth has been slow during this recovery. Second, the output gap is notoriously difficult to estimate because potential output is difficult to estimate. Our paper, in fact, seeks to understand and quantify reasons why potential output growth has been so weak. Quantifying that requires adjusting for the cycle, and the view in this paper is that the best estimate of the cycle is obtained using Okun's Law.

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.

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

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

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 (Robinson 1988, Stock 1989; for technical conditions on μ and recent references see Zhang and Wu (2012)). In this approach, μ is estimated as a long-run smoothed value of y, after subtracting off 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.[3] 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.

[3] An alternative approach for identifying and estimating μ is 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 sounddifferent, 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 give quite similar results. However, because the state-space approach entails estimation of different model parameters for each series,

Estimation details. 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 $p = q = 2$ leads and lags judgmentally 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 results are robust to using more leads and lags. We compute heteroskedasticity- and autocorrelation-robust standard errors using standard formulas for the sampling variance of predicted values.

The use of this regression departs from the norm in partially linear regression, in which β is estimated by regressing prefiltered $(1-\kappa(L))y_t$ on leads and lags of prefiltered $(1-\kappa(L))\Delta u_t$, however 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.

For the low-pass filter, we use a biweight filter with truncation parameter 60.[4] 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 productivity.[5]

---

the implied smoothing filters differ across series so the additivity property discussed in the next subsection does not hold for the state space approach.

[4] Tukey's biweight filter w(L) is two-sided with $w_j = c(1 - (j/B)^2)^2$ for $|j| \le B$ and = 0 otherwise, where B is the bandwidth and c is a normalization constant such that $w(1) = 1$.

[5] 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.

***Additivity*.** The foregoing method for estimating the trend, cycle, and irregular parts has the useful property that it preserves additivity when applied to the growth accounting (or any other) decomposition. Specifically, suppose that $y_t = y_{1t} + y_{2t}$. Then this additivity is exactly 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 κ(L) for all series, and the property that regression is linear in the dependent variable.

II.B Method 2: Using a dynamic factor model to estimate the cycle

The dynamic factor model produces forecasts of the variables under study using history on a broad cross section of macro variables through 2009, second quarter. Specifically, 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 have 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 completeness; for additional detail, including variable transformations and measures of fit, see the online appendix to this paper.

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

$$X_t = \Lambda F_t + e_t \tag{10}$$

$$F_t = \Phi(\mathrm{L})F_t + \eta_t, \tag{11}$$

where Λ is the $N\times r$ matrix of factor loadings, Φ(L) is are the vector autoregression (VAR) lag polynomial for the factors, X_t is the $N\times1$ vector of series, F_t is the $r\times1$ 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 a detailed description is available in the supplemental data appendix.[6] The forecasts are computed as follows.

(i) All 123 series used to estimate the factors are transformed to approximate stationarity.[7]

(ii) The factors are estimated by principal components (computing using least squares on the unbalanced panel of data) over the period 1959q3-2016q2.

(iii) The DFM parameters Λ and $\Phi(L)$, with 4 factors and 4 VAR lags, are estimated by OLS, using data from 1984q1 to 2009q2[8], treating the estimated factors $\hat{F}_t$ as data.

(iv) Given the factors through 2009q2, forecasts of the factors, $\hat{F}_{t|2009q2}$, are computed for t = 2009q3-2016q2 using the factor VAR and history of the factors through 2009q2.

(v) Given the factor forecasts, forecasts of the detrended variables are computed as $X_{t|2009q2} = \hat{\Lambda}\hat{F}_{t|2009q2}$ for t = 2009q3-2016q2, where the estimated value of Λ is computed by using data from 1984q2-2009q2.

---

[6] 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.

[7] 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. Any remaining near-zero frequency variation is removed by local demeaning using a biweight kernel with 25-year bandwidth. See Stock and Watson (2016).

[8] The start date of 1984q1 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 in Stock and Watson (2016), 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.

(vi) Forecasts for the original series (not detrended) are computed by adding the forecast of the detrended variable to the 2009q2 value of the trend, adjusted as appropriate for the demographic trend in the labor force participation rate (details discussed below).[9]

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.

Post-2009q2 trend for forecasts. To maintain the pseudo out-of-sample approach used here, the post-2009q2 trends freeze the trends in each series at its 2009q2 value (with one exception), and project forward 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 et al. (2006) and CBO (2007). Here, we adopt CBO's (2007) method for projecting the decline of participation, which is to freeze age-gender participation rates at a fixed date and to compute overall participation rates as the weighted average of age-gender participation rates times evolving age-gender population weights.[10] This 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 CPS employment, and hours per employee unchanged.) The result is a trend in real GDP that incorporates aging effect on employment as understood in 2009q2, with other component trend growth rates frozen at their 2009q2 values. Trend growth

[9] The results are robust to variations in the benchmark model, including shifting the jumping-off date to 2009q4. For the main series, including GDP and employment, and productivity, they are also robust to using a low-dimensional VAR.

[10] Specifically, using CPS microdata, we estimate participation rates for men and women separately at each age 16-80. These participation rates are estimated using 2005 and 2006 data to reflect participation at full employment (during the previous expansion, the unemployment rate first fell below the CBO estimate of the NAIRU in 2005q3). We then use actual CPS population shares to compute counterfactual changes in participation that would arise solely from demographic shifts (CPS age-sex population shares are extended into 2016 using scaled Social Security Administration share projections).

rates of the demand components of GDP are computed as the component's time series trend as of 2009q2, plus the share-weighted difference between the GDP trend (inclusive of the participation aging trend) and the 2009q2 value of the GDP trend. This adjustment ensures that the share-weighted trend growth rates add.[11]

***Additivity and share weighting*.** Like the Okun's Law method, this 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 components in 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 4 to provide standard units of percent change in output per percentage point change in the unemployment rate. The second is the standard deviation of the cyclical part with, for comparison, the standard deviations of the trend and irregular parts. The third, in the final column, is the R^2 of the common component of the variable in the DFM, that is, the R^2 in equation (10) of the variable against the estimated factors. Lines 1-5 present results for various measures of aggregate output. Lines 6-14 provide results for real business output per capita and the variables entering its per-capita growth decomposition (2).

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

[11] This final adjustment is inconsequential because the 2009q2 participation adjustment to the trend value of GDP is small.

Okun's law coefficient on the sum of the components; that is, the coefficients in lines 7-10 add to -2.02, the coefficient on real business output per capita. Thus 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 in business output per capita, as measured by the generalized Okun's law coefficient of -2.02, slightly over one-third (-0.71) comes from the employment rate (one minus the unemployment rate), one-eighth (-0.23) comes from variations in hours per worker, and a small amount (-0.11) comes from the labor-force participation rate. These results support the general view that participation is slightly procyclical, falling as unemployment rises. Of course, a large unexpected reduction in participation occurred after the 2008 financial crisis. We take up the subject of how much of the recent decline in participation can be related to the slack labor market later in this paper (Section V).

Interestingly, one-fourth (-0.47) comes from cyclical variation in the ratio of business employment to CPS 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 CPS. Fernald and Wang (2016) find that hours worked has almost the same cyclicality in the two surveys. Together, cyclical variation in hours per capita accounts for three-fourths (-1.52) of the cyclical variation in business output per capita.

The remaining cyclical variation in business output per capita is accounted for by substantial cyclical variation in productivity (-0.50). Research on productivity 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), and variations in productivity account for one-fourth (-0.50) of the cyclical variation in business output.

Finally, cyclical variation in capital input per capita and labor quality is much smaller than for the other variables. Though investment is pro-cyclical, the cumulated stock of capital changes relatively little in synchrony with unemployment. The positive coefficient on labor quality (0.09) supports the hypothesis that times of high unemployment are times of higher labor quality, because unemployment differentially strikes lower-skill workers.

The remaining columns quantify the amount of variation in the variable that is cyclical, as measured by first by the variation in the Okun's law estimate of c_t and second by the fraction of the variance of the series explained by the factors. 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 productivity accounts for one-fourth of the cyclical variation in business output per capita, cyclical variation only accounts for a fraction of the variation in productivity. This finding indicates a large amount of high-frequency variation, including measurement noise, in productivity growth.

III.B Method I: Using Okun's law to estimate the cycle

Table 3 summarizes the results of the growth accounting decomposition, in which the cyclical component is estimated using Okun's Law and is conditional on the unemployment rate path. As a point of reference, Table 3 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 2009q3 to 2016q2) or the trough-to-peak period, whichever is shorter.

The left panel of three columns in Table 3 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.

Figure 4 through Figure 12 plot the trend, cyclical, and irregular parts of business output per capita and its growth-accounting components. The organization of the figures is the same for each series, so this explanation focuses on figures for business output per capita. Figure 4 shows the rate of growth of the low-frequency non-cyclical part of our measure of output, real gross domestic output per capita. The horizontal line is the average annual growth rate over our sample period. The trend of growth started just above average in the early 1980s, rose to a maximum in the late 1980s, then dropped to low levels in the 2000s and the current decade. Figure 5 shows all three parts. The cyclical and irregular parts are cumulated from the log-first differences that emerge from our estimation. If the trend part were cumulated, it would be difficult to distinguish from a straight line and would be on a different scale. To focus on variations in the trend, the figure shows the cumulative gap between the trend and its average, that is, the integral of the area between the trend and the mean in Figure 4. Thus, the low-frequency non-cyclical part reaches a maximum in the second graph at the time when the line in the first graph cuts the horizontal line showing the mean of the estimated first difference.

III.C Method 2: Using the dynamic factor model to extract the cycle

Table 4 presents 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 2009. The first column, the forecast, is the sum of the cyclical component of the forecast and the trend, averaged over 2009-2016. 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 (10)) is given in the final column.[12]

[12] Note that that 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.

Figure 13 plots the predicted and actual paths of the growth accounting decomposition in Table 4. The gap between the predicted and actual is the irregular part z_t, computed using the DFM, conditional on the state of the economy in 2009.

III.D Discussion

An important difference in the results for the 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 very 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 in 2009q2, and in principal the actual path of any variable, including labor market variables, can depart arbitrarily from its forecast path. In fact, the factor forecasts under-predict the robust recovery in the labor market, but over-predict the growth of GDP. This robust recovery of employment combined with the slow growth in output is a key feature of this recovery, and we return to this and other implications of the factor forecasts in more detail at the end of this next section.

Aside from this important difference in forecasts of the labor market 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 recent recovery. For clarity, we therefore focus primarily on results using the Okun's law 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. In the previous three recoveries, GDO growth averaged 3.57 percent per year, while in the current recovery, it averaged 2.20 percent per year, for a shortfall of 1.37 percentage points. Similarly, business output per capital averaged 2.48 percent per year in the current recovery, compared to 1.02 percent in the previous three, for a shortfall of 1.45 percentage points. The decomposition of this shortfall in column 3 splits this decline in business output into three factors, which make nearly equal contributions: a decline in the growth of capital per capita (capital shallowing), a decline in the growth rate in productivity, and a decline in the participation rate. The other contributions are smaller and largely offsetting.

This comparison of actual growth rates understates the shortfall in output, however, for it does not take into account that the recent recession was much deeper than were the previous three on average. The second block of estimates therefore presents this decomposition after subtracting the cyclical component of these series, as estimated by Okun's law, conditional on the unemployment rate.

Making this cyclical adjustment creates a different, starker picture of the slow growth. The shortfall in GDO is now much larger, 1.81 percentage points. Cumulatively over this recovery, the shortfall in GDO is nearly 14 percent (right-hand column). The shortfall in business output per capita is also larger after cyclical adjustment, also 1.81 percentage points. After cyclical adjustment, the single largest source of this shortfall among the growth accounting components arises from slow productivity growth, accounting for 0.71 percentage points—more than one-third—of the shortfall in business output. Shortfalls in the participation rate and in capital input per capita are also large and have similar contributions, respectively 0.54 and 0.53 percentage points.

Business output. Figure 5, 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 non-cyclical 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.

***Capital input*.** Figure 8 shows the movements of the level of weighted capital input, using the same scale as for output and productivity. Capital input contributes moderate amount of non-cyclical movement to output, even though 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 productivity, the labor force, employment, and other endogenous variables. The movements of capital input per capita shown in the figure reflect the joint determination of the variables. For example, if the economy is hit by an exogenous decline in the rate of growth of productivity, optimal capital input grows less rapidly as well, according to most models of investment. In a different decomposition of the movements of output that stated inputs as ratios to output (based on equation (5)), the shortfall of capital per unit of output after the crisis would essentially disappear.

***Hours per worker*.** Figure 9 shows the levels of the three statistical parts of weekly hours per worker. Consistent with the coefficient of -0.23 in Table 2, the cyclical part of hours rose smoothly during the recovery, as in the three earlier recoveries. The low-frequency trend also rose slightly, while the high-frequency non-cyclical 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 10 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. Figure 11 shows the levels in the same format as earlier graphs of levels, with the low-frequency part stated around mean shrinkage. 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.

Line 10 shows that the bridging needed to unite standard labor-market measures such as unemployment with standard productivity research referring to the private economy has little effect for the cyclical part but quite a bit for the non-cyclical part.

Finally, Figure 12 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.

***Comparison to the factor-model estimates of the cycle and irregular parts*.** Table 4 shows that compared to what would have been expected based on the data through 2009, actual GDP growth fell short by 0.51 percentage points, GDO growth by 0.37 percentage points, and business output by 0.29 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 in Table 3 conditions on the unemployment rate path. For example, the contribution of actual average establishment employment growth exceeded the forecast path by 0.34pp. Relative to the factor-model forecasts the recovery in the labor market was stronger than expected, while the recovery in output was weaker than expected. 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.

The results in Table 4 confirm the main conclusions of the growth accounting comparison in Table 3 based on Okun’s law. Specifically, the largest single contribution to the shortfall in output is the shortfall in productivity growth, with a forecast error of -0.52 percentage points. Other large negative forecast errors were associated with participation (-0.19 percentage points) and capital input per capita (-0.19 percentage points). These contributions of productivity growth, participation, and capital input, are quantitatively close to the corresponding contributions to the irregular part of output growth using the Okun's laws method, which respectively are -0.35 percentage points, -0.19 percentage points, and -0.12 percentage points (the difference between the two productivity contributions is in part due to the different treatment of trends). Despite the slower than expected growth the participation rate, per capita labor hours grew faster than forecast (making a -0.51 pp contribution) because of stronger than expected growth in hours per worker, business employment relative to household-survey employment, and the employment rate. We return to these labor market factors in Section V.

IV Why Have Capital Accumulation and Productivity Fallen Short?

The previous sections found that shortfalls in productivity growth, labor-force participation, and capital accumulation account jointly most of the shortfall of output growth relative to the previous three recoveries—see Table 3. This section first argues that the capital shortfall is mainly endogenous, given the shortfall in productivity and labor input. We are not staking out a strong position on this point and there is no assumption about endogeneity or exogeneity in the calculations underlying that table. Rather, we observe that research has not yet uncovered any compelling mechanism linking productivity growth or participation to other economic forces in the past decade, whereas those two forces rather easily account for the shortfall in capital input.

The first part of this section lays out the case that the behavior of capital formation is not hard to understand conditional on the shortfalls of productivity and labor input. The second part examines evidence and hypotheses regarding the productivity shortfall. Much of that shortfall appears structural, reflecting pre-recession trends. Our preferred hypothesis is that it 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.

IV.A Why has capital fallen short?

Many commentators have emphasized the slow pace of investment during the recovery. On its face, this concern seems appropriate. Figure 15 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 tight credit for some borrowers, heightened uncertainty, regulatory barriers, or other special factors during the recovery. 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, the labor-productivity growth accounting in equation (5) should show large deviations from previous experience. Table 5 presents calculations based on that equation that assign only a small role to a capital shortfall (Table 6 presents the analogous calculations, based on the forecast-based method of cyclical adjustment). The top panel shows that business-sector labor productivity growth (output per hour) was more than a percentage point per year slower in the post-crisis recovery compared with the three previous recoveries. A shortfall of growth in productivity and the capital-labor ratio explains almost all of that shortfall.

The bottom panel uses the second version of the identity in equation (5). The capital-output ratio contributes essentially nothing to the labor-productivity shortfall. In other words, capital growth has been weak, and labor growth has been weak, but the ratio of the two has been normal. This is consistent with the view that weak productivity and labor explain the shortfall.

Another way to view this is through the return to capital. An important determinant of business investment is the payoff to owners of capital. Some stories for weak investment imply that perhaps capital was not earning as much as in normal times. But, as Figure 16 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 remained normal in the succeeding years of the recovery.

IV.B When did productivity growth slow?

TFP growth varies over time. The view in this paper is that the fundamental determinants of growth have an important structural, non-cyclical component. In terms of TFP, many macroeconomic models assume that trend growth is constant over time. But historical and statistical analysis suggests otherwise. For example, Gordon (2016) documents that U.S. TFP growth picked up after the late 19th century, was then relatively strong from 1920-1970, and has mostly (apart from the 1995-2005 period) been slower since then. He argues that some major inventions are more important for economy-wide innovation than others. Alexopoulos and Cohen (2011) document lower-frequency waves of innovation as measured by time-series variation in publications of technical books over the 20th century. Even on more recent periods, break tests on TFP or labor productivity growth suggest that trend growth varies over time (e.g., Fernald, 2007; see also below). Kahn and Rich (2007) argue that a regime-switching model with high- and low-growth periods describes U.S. productivity in the post-war period. And some growth models suggest that U.S. growth (per capita or per hour) has been above its steady-state pace since the 19th century—again implying growth may change (Jones, 2002, Fernald and Jones 2014).

Professional forecasters certainly view underlying trend productivity growth as varying. Figure 17 plots the median forecasts from the Survey of Professional Forecasters for growth in non-farm business output per hour over the next 10 years. (The surveys start in 1992 and are done in the first quarter of the year.) The 10-year-ahead forecasts range from 1.3 percent to 2.5 percent. They broadly track the lagging 10-year average of actual productivity growth. They rise sharply between the 1999 and 2000 surveys, then for the most part remain close to 2.5 percent through the 2006 survey. Between 2006Q1 and 2008Q1 (before the intensity of the financial crisis was felt), the surveys fell to 2 percent. Hence, even before the recession, forecasts had come down markedly.[13] The persistent weak productivity growth ever since has led to further reductions in forecasts.

Given this evidence for time variation, it is no surprise that Müller and Watson (2015) find, statistically, that confidence intervals on average U.S. productivity growth over long periods of time (say, over the next 10 or 25 years) are very wide.

Timing of the TFP growth slowdown. The visual evidence above suggests that trend TFP growth slowed prior to the recession. For example, the noncyclical components of TFP discussed in Section III (the low-frequency trend and the irregular component) both peaked in the early- to mid-2000s (Figure 6 and Figure 7). Of course, that decomposition assumed a particular value for the biweight filter (with window of 60 quarters). Figure 18 shows the underlying data on the level of TFP, unemployment-adjusted TFP, and the biweight-filtered trend. It also shows a broken linear trend line, described later (Levels are calculated by cumulating 100 times the quarterly log changes since 1981:Q2.)

[13] The lagged 10-year data are not yet the real-time data that forecasters saw at the time. John's RA is working on that. In addition, Fernald (2015) documents that every data revision from 2005 to 2009 reduced the growth rate of productivity after 2004—making the pre-recession slowdown more apparent ex post than forecasters saw in real time.

Visually, it is clear that the growth rate of both TFP and cyclically-adjusted (i.e., Okun-adjusted) TFP rose in the mid-1990s but flattened out after the mid-2000s. TFP has a sharp V-shaped decline and rebound during/immediately after the recession. The Okun unemployment adjustment flattens out the V shape, but the low-frequency changes in trend are apparent in both series.

Break tests confirm this visual impression that the TFP trend appeared to change prior to the recession. Focusing on results for cyclically-adjusted TFP, Bai-Perron tests for multiple structural change in the mean of the growth rate show evidence of two breaks over the 1981Q2-2016Q2 period. First, the tests reject the null at the 5 percent level of no breaks against the alternative of an unknown number of breaks (the so-called UDMax and WDMax tests). Second, two breaks are significant at the 5 percent level. The point estimates are that TFP growth rose after 1995Q4 (check) and slowed again after 2006Q1. Finally, break tests on labor productivity, which incorporates the endogenous capital deepening response, show an even stronger, and earlier, slowdown.[14]

The figure also shows a broken linear trend corresponding to the estimated break dates. This line is visually helpful for focusing attention on the changes in trend, and highlights how slowly the biweight-filtered trend adjusts. (Of course, the broken linear trend is somewhat misleading since the break tests assumed that TFP growth was a random walk with potential breaks in the drift term; and (preliminary) statistical tests point to the latter characterization of the time series process.)[15]

Fernald (2015) discusses other evidence of a pre-recession slowdown in growth, especially from industry data. For example, the pre-recession slowdown in TFP growth was

[14] Break tests are preliminary. The exact date of the mid-2000s date varies based on series. For TFP or unemployment-adjusted TFP, the break is dated 2006Q1, or 2005 in annual data. For Fernald's utilization-adjusted TFP measure, the point estimate is 2004Q4 (annual 2005). For annual labor productivity data, the break is after 2004.

[15] [Details TBD]

broad-based across industries—it was not, for example, particularly pronounced in housing-related industries or finance.

The discussion above focuses on outcomes—i.e., 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 TFP.

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 14 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.[16]

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). Indeed, in U.S. data

[16] 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.

since the early 1970s, the unusual period for TFP growth was roughly the decade from late 1995 to 2005.

We now turn to hypotheses for the slowdown in TFP.

IV.C Explanations for slow TFP growth

This section surveys the main explanations that have been proposed for the slow pace of TFP growth since the mid-2000s. We put the most weight on it being a pause in (if not an end to) the information-technology (IT) revolution. The reason is that it fits a body of empirical and theoretical evidence; plus, other explanations have shortcomings. The explanations we consider, with a short summary of the takeaway for each, are the following:

- Mismeasurement. Existing evidence does not suggest that mismeasurement of business-sector growth has grown worse since the mid-2000s. Moreover, the scale of potentially omitted Internet-based products is much too small to account for the shortfall in TFP growth. Thus, the slowdown in the pace of TFP growth since the mid-2000s appears real.
- Fallout from the recession and financial crisis: Recessions can sometimes leave "long shadows" on growth and the level of productivity (Fatas, 2000)…but don't always. For example, the Great Depression was a highly innovative period and, in advanced economies, there is less evidence than for emerging markets that recessions (financial or otherwise) leave a permanent imprint on TFP.
- Rising regulation/loss of dynamism. Measures of dynamism have been declining for decades. One hypothesis for this decline is rising regulation. The time series timing does not appear to line up, and the existing evidence does not yet indicate to us that this is a quantitatively important factor in explaining slow recent growth.
- A pause in the IT revolution. Considerable research links the mid-1990s pickup in TFP growth to the production and use of IT. This story has the virtue that it fits the anecdotes, the time-series data, and a considerable body of economic research.

The latter three explanations all presume that the slow growth in TFP is a real phenomenon. But the different diagnoses do point to different policy recommendations. We now discuss the alternative hypotheses in greater detail.

1. Mismeasurement. Perhaps the problem of slow growth in both productivity and GDP is illusory? That is, perhaps we aren't fully measuring the gains from IT-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 IT-related growth has gotten *worse* since the early 2000s. For example, Byrne et al. find that prices for computers and mobile phones were mismeasured in the 1990s—when domestic production of these products was a much higher share of GDP than it is today. So for GDP and productivity, this source of mismeasurement turns out to be quantitatively more important in the 1990s and early 2000s than more recently. More broadly, 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 TFP growth was broadbased across industries. So changes in weighting matter relatively little.[17]

Finally, the great benefits we all get from smartphones and from free digital services such as Facebook and Google certainly improve leisure and home production. But while valuable, these benefits are largely outside the market sector. They do not change the fact that business-sector TFP growth has been slow.

2. <u>Fallout from the recession and financial crisis.</u> As shown in Table 2 and as is evident in Figure 7 and Figure 13(d), productivity has a large cyclical component, which accounts for the surge in productivity in productivity growth in the first few quarters of the recovery and for some of the slow growth in productivity since 2013. Because the recession was unusually large, these cyclical movements in productivity growth were unusually large. But our accounting adjusts for these normal cyclical movements, so any explanation of additional slow productivity growth

[17] More broadly across the economy, Aghion et al. (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.

arising from this recession must relate to special features of the financial crises, or perhaps nonlinearities because of the depth of the recession. On this latter point, theory is ambiguous about whether severe recessions (including financial ones) have a persistent effect on the path of TFP (its growth rate or just its level). Intuitively, a crisis could reduce the invention or adoption of new technologies (e.g., 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. Empirically, 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 when the economy recovers. If weak TFP growth is 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).

That said, the reallocation effects in some models go the other way, raising measured TFP in a credit crisis (e.g., Petrosky-Nadeu, 2013, or the "cleansing effects" of Caballero and Hammour, 1994). Bloom (2013) points out that high uncertainty can even stimulate longer-run innovation.

Overall, there is limited empirical evidence for *developed* countries that previous business-cycle downturns (financially related or otherwise) permanently harm the level or growth rate of TFP. 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) look at growth-accounting variables across countries following financial crises. For developed countries (but not for others), the long-run level of TFP is essentially unchanged by a financial crisis—indeed, the point estimate is a small positive. For the United States, Huang, Luo, and Starts (2016) find that the level of TFP always bounces back quickly from recessions, including after 2009.[18]

The biggest challenge for explaining U.S. data is the timing. TFP growth appeared to slow prior to the recession. Anzoategui et al. (2016) argue that there was a pre-recession shock to exogenous growth combined with the large shock from the recession. Still, as noted in Section IV, there is limited U.S. evidence that investments in R&D and other intellectual property slowed because of the recession. Rather, the pace of growth slowed earlier.

3. Rising regulation/loss of dynamism. These explanations are conceptually distinct, though related. By many measures, the dynamism of the U.S. economy has declined since the 1980s (e.g., Decker et al., 2016a, b). The reasons are not entirely clear, but regulation is one story (Barro, 2016). For an example of how regulatory barriers might matter, Cette, Fernald, and Mojon (2016) compare the gains from IT in the U.S. and Europe. Europe invested in computers, but didn't get the same productivity benefits after 1995. The leading hypothesis is labor and product-market inflexibilities—many induced by regulations—that have limited the ability of firms in Europe to reorganize to benefit from IT.

In the U.S. context, however, there is less evidence that this is the primary issue. First, a challenge is timing. In particular, some commentators focus on regulatory changes since 2008. But as already noted, relatively slow productivity growth has been the norm since the 1970s (with the exception of the decade after 1995). Moreover, the downward trend in dynamism is longstanding; Decker et al. (2016b) suggest that the character of declining dynamism changed after 2000 or so, rather than being a post-2008 phenomenon. Of course, that earlier date would match the view that there were structural shifts in trend growth prior to the 2007-09 recession. Second, Goldschlag and Tabarrok (2014) find empirically that changes in U.S. Federal

[18] Some recent/upcoming IMF work suggests more of an effect [to be added]. Still, 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 (e.g., Cette, Fernald, and Mojon, 2016).

regulations have little or no effect on industry entrepreneurial activity or dynamism. Third, causation between innovation and dynamism is also not necessarily clearcut. For example, if the "problem" is the lack of available/exploitable ideas (low-hanging fruit) for the broad economy, then the lack of dynamism might simply be a symptom of that lack of opportunity.[19]

In any case, whatever the flow of new ideas turns out to be, business flexibility is likely to be needed to take advantage of it. And there is a bipartisan view that many regulations—such as overly restrictive land-use restrictions and onerous occupational licensing—constrains activity. Still, it is unclear that this is the first-order issue explaining recent slow TFP growth.

4. A pause in the information technology revolution. The hypothesis that IT was the culprit is natural. A large literature links the mid-1990s speedup in TFP growth to the exceptional contribution of computers, communications equipment, software, and the Internet. The idea is that IT has had a broad-based and pervasive effect on the economy through its role as a "general purpose technology" (e.g., 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. As a result, businesses throughout the economy transformed how they operated and became more efficient. But, by the early 2000s, industries like retailing had, plausibly, already been substantially reorganized, after which the gains from further innovation might have been more incremental than transformative.

Fernald (2015) finds that, in industry data, the slowdown after the early 2000s (prior to the Great Recession) is mainly accounted for by industries that produce information technology (IT) or that use IT intensively. We might see another such period in the future, perhaps reflecting

[19] Some anti-regulation advocacy groups point to Federal Register pages. In a growing economy, it is not surprising that the page count grows over time. But the growth rate has been slow since 2000 or even since 2008 (Using data from http://sdrv.ms/1740Gv1). Relative to GDP, the number of pages in 2016 was at its lowest level since 1971.

cloud computing, the Internet of things, and the radical increase in mobility from smartphones. However, we have not yet seen those gains in the data.

This story has a number of virtues. First, it is consistent with the large literature on the role of IT in the productivity acceleration in the late 1990s. Second, it is consistent with the view in the GPT literature that the gains are, essentially, a series of drawn-out levels effects. Ex ante, it is hard to predict how long the gains will continue. Indeed, the gains might ebb and flow for a time (e.g., Syverson, 2013). Finally, the story matches the facts in Figure 18 as well as the industry data adduced by Fernald (2015).

To conclude this section, it is difficult to know the counterfactual for how TFP growth would have evolved without the recession, or with a different regulatory regime. But the evidence suggests that the slow pace is real, and that it is a continuation of pre-recession trends.

V Changes in the Labor market

Table 3 shows that a decline of participation in the labor force accounted for a shortfall of about 3.7 percent in real output per capita during the post-crisis recovery, after adjustment for the normal cyclical rise when unemployment falls. Figure 11 shows that the shortfall was in both the irregular and trend parts. 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 -0.11 over a sample period that includes the rise in unemployment and fall in participation during and after the crisis recession. The coefficient is even closer to zero for samples ending before the crisis.

The labor force comprises people 16 and over who are working or are actively looking for work. Trends in participation have been quite different for men and women, so it is a good idea to consider them separately. Figure 19 shows the percentages of the populations who counted as in the labor force starting in 1984. The rate for men declined smoothly prior to the crisis, then faster in the post-crisis years, and finally leveled out in 2015 and 2016. The rate for women grew until 1996, continuing a low period of growth in earlier years, flattened until 2008, declined noticeably during the recovery, and leveled out in the last two years.

Figure 19 shows that participation by both men and women fell noticeably—by about three percentage points—after 2008. Answering the counterfactual—what would have happened to participation had the trauma of 2008 and the long slump following not occurred?—is a challenge. But it seems likely that some forces specific to the post-crisis years depressed participation.

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 adjustment 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 rather more likely to finish high school and attend college; thus 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.

Figure 20 illustrates the effect of controlling for education in the computation of aging trends in the participation rate. Adjusting for age and sex alone explains 2.0 pp of the 3.4 pp drop between 2006 and 2015, however controlling in addition for education reduces this demographic component of the decline to 1.3 pp (although controlling additional for race and marital status raises this to 1.9 pp). While the aging baby boom undoubtedly is and will be retiring, the timing of that retirement – and thus the contribution of this aging effect – to the decline in the participation rate over the recovery is not clear cut.

Figure 21 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* by 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 7 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.

An important fact for interpreting the decline in participation during the recovery is that the joint movements of participation and unemployment from 2008 to 2016 have been almost exactly uncorrelated. There is no support for the hypothesis that participation fell on account of a slack labor market based on the traditional criterion of the contemporaneous regression coefficient. Rather, participation fell as unemployment rose, from 2008 to 2010, and then participation continued to fall as unemployment returned to normal. The regression coefficient on the same basis as the one in Table 2, but for the recovery period alone, is -0.01 with a standard error of 0.11.

Erceg and Levin (2014) studied state-level data on participation and unemployment from the BLS's Local Area Unemployment Statistics program. The main source of the data is the Current Population Survey. 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 with a standard error of 0.08, in their preferred specification (their Table 2, p. 12). They conclude that “These regression results provide stark evidence that cyclical factors have been crucial in explaining the recent decline in prime-age participation” (p. 11). We believe that their interpretation is an unjustified leap from correlation to causation. Their evidence is strong that states with higher post-crisis increases in unemployment also experienced larger declines in participation. The finding is that changes at the state level mimicked the national change but were not as pronounced: The national unemployment rate rose by 5 percentage points in annual data from 2007 to 2010 and the participation rate fell by 2.3 percentage points, for a ratio of -0.46, well above the coefficient that Erceg and Levin found.

The negative relationship between unemployment and participation in state unemployment data for the post-crisis period is not typical of the data for earlier periods. Over the entire duration of the state data, from 1976 to 2016, the coefficient for monthly changes of participation conditional on monthly changes in unemployment is +0.1103 with a standard error of 0.00004, the opposite of the correlation Erceg and Levin found for the post-crisis period.

The surprising, large, and persistent decline in labor-force participation is a phenomenon that deserves and will receive intensive study. The simple idea that it was a response to an extremely slack labor market in the years immediately following the crisis has 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 of Slow Output Growth

So far, our discussion has focused on understanding this 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 the weak growth. We therefore now turn to some of those other explanations.

Before considering these ideas individually, we note that our earlier results based on the use of unemployment as a cyclical indicator, and on a factor model with a multivariate statistical characterization of the cycle, take demand into consideration. For example, 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 weak demand are ruled out. Sponsors of explanations based on weak demand need to couple their explanations with a parallel explanation of why labor-market indicators misleadingly pointed to full employment by 2016 when actually there was substantial excess supply of labor that could accommodate higher output levels.

We group the remaining explanations into three broad categories. The first group emphasizes chronic problems related to inadequate demand or policy failures: secular stagnation; retarded consumption growth because of growing inequality; hysteresis that manifests as slow growth of the labor force and employment; and failures of monetary or fiscal policy. The second group of explanations comprises one-off explanations: weak international demand as Europe struggled to recover and growth slowed in China; lingering effects of the financial crisis on both the housing sector and private fixed investment; fiscal headwinds because of the combined unwinding of fiscal stimulus from the American Reinvestment and Recovery Act and the sequester; fiscal policy uncertainty that postponed investment and consumption and thus depressed demand; and the potentially reduced ability of monetary policy to provide effective stimulus in the face of the zero lower bound on interest rates. The third group of explanations suggests that at least some of the slowdown in output growth (and in productivity) is not real but rather 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

To shed light on these explanations, we use the rich set of internally consistent forecast errors generated from the forecasting exercise. To this end, Table 8 presents a contributions-accounting decomposition of the forecast of output and its main components, based on the demand-side national-income accounting identity. The entries in this table are contributions to mean growth, computed using share weighting; for example, the entries in the first column correspond directly to Table 2 in the GDP release ("Contributions to Percent Change in Real Gross Domestic Product"), which presents contributions of components to growth rates, except that here these contributions are averaged over 2009, third quarter, through 2016, second quarter. 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. Figure 22 plots the predicted and actual paths of selected NIPA variables in Table 8 (also not share-weighted), and Figure 23 plots predicted and actual paths of employment growth.

Examination of Figure 22 and Figure 23 indicates three periods in the history of this expansion. From 2009q3-2010q4, the economy grew strongly, with employment, GDP, consumption, and private fixed investment all growing at or above the forecast path. From 2011until2013, although employment growth was strong, it was below its predicted path, and the associated predicted strong growth in output failed to materialize. This is the period of the large growth gap – the notable absence of sustained output growth in the 3 to 4 percent range. After an initial surge in 2009, the growth of productivity was low during this period, well below the predicted path (Figure 13(b)). 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. Thus the picture is one of a recovery delayed: the slow-growth puzzle is largely the absence of strong growth in GDP and in productivity during 2011 through 2013.

The demand decomposition in Table 8 indicates that most of the demand components tracked, on average, their forecast paths. For example, although exports were unexpectedly weak, so were imports, after share-weighting their contributions to the average forecast error in output growth is negligible, -0.02 and -0.01 percentage points respectively. Table 8 indicates that the average forecast error is largely attributable to three source: consumption of services (-0.13 percentage points), federal government expenditures (-0.19), and state and local government expenditures (-0.11).

For federal government spending, the main shortfall occurs in 2013 and 2014. This period coincides with the fiscal drag associated with Recovery Act expenditures were unwinding and the sequester being put in place. 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, and made a modest -0.19 percentage point contribution to output growth. Most of this weakness is attributable to two service sectors: housing and utilities (-0.6 percentage points) and financial services and insurance (-0.6 percentage points). The shortfall in financial services and insurance is largely associated with a large negative value in 2012 that appeared in a recent revision of the series.

The forecast error in residential investment averaged 0.09 pp 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 positive average contribution of residential investment to the GDP forecast error.

VI.B Discussion

Because we do not identify a structural factor model, we do not identify the structural shocks that led to the 0.51 percentage points of slow GDP growth over this period. Nevertheless, these patterns of forecast errors shed 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 unpredicted shortfall in GDP growth was only -0.19 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 22(h)).

The fact that, together, 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 since employment growth exceeded the 2009 prediction on average, due to strong growth 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.

This examination of the components of demand does support one contribution to the slow recovery: the weakness in both federal and state and local government spending. 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 0.19 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 expenditures 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 spending associated with balanced budget requirements and the prolonged effect on real estate tax receipts of the fall in house prices during the recession.

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 8 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 GDP 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 Remark

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, though both phenomena have so far escaped anything close to a definitive explanation.

Data Appendix

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.

The Fernald dataset further decomposes standard productivity growth (defined implicitly by equation (2)) into growth in factor utilization (based on a model and estimates from Basu, Fernald, and Kimball, 2006) and in "utilization adjusted productivity." We do not use that decomposition in this paper. The reason is that we undertake the same cyclical adjustment of all variables, including productivity, by projecting them on unemployment-rate changes. Because our decompositions are additive, we could use the Fernald decomposition of standard productivity. We confirmed that our conclusions would not be affected, since the cyclical adjustment based on the unemployment rate leads to a qualitatively similar adjustment as the Fernald utilization-measure. In particular, the productivity shortfall we document in the paper is still large. For simplicity, and to avoid relying on a particular economic model of cyclical adjustment, we focus on the standard measure of productivity.

For the overall economy, and for the business sector, we combine expenditure-side and income-side measures of output. The "standard" measures of aggregate output (gross domestic product, or GDP) and business-sector output from the Bureau of Economic Analysis (BEA) are from the expenditure side. However, Nalewaik (2010) suggests that gross domestic income (GDI) may better capture the business cycle variations in output growth. He finds that GDI correlates more strongly with other business cycle variables, before and after data revisions. Nevertheless, Greenaway-McGrevy (2011) and Aruoba et al (2012) suggest that both GDP and GDI provide independent information, and recommend taking a weighted average of the two. We use the geometric average of the income and expenditure sides, which (following Council of Economic Advisers, 2015) we label gross domestic output (GDO).

The standard measure of business output in the NIPAs and in the BLS productivity releases comes from the expenditure side. In particular, it is GDP less non-business output. Fernald (2012b) constructs a corresponding income-side measure as GDI less non-business output. Our benchmark measures of business-sector output, labor productivity, and productivity use an equally weighted average between the GDP and GDI based measures of business output.

References (incomplete)

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* Spring 2006: 69-154.

Aaronson, S., T. Cajner, B. Fallick, F. Galbis-Reig, C. Smith, and W. Wascher. 2014. "Labor Force Participation: Recent Developments and Future Prospects," *Brookings Papers on Economic Activity*, Fall 2014, 197-275 (including discussion).

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

Alexopoulos, Michelle and Jon Cohen. 2009. "Volumes of evidence: examining technical change in the last century through a new lens." *The Canadian Journal of Economics* Vol. 44, No.2 (May 2011), pp. 413-450.

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, 2016. http://scholar.harvard.edu/files/barro/files/wsj_published_version_092116.pd f

Susanto Basu and John G. Fernald, 2001. "Why Is Productivity Procyclical? Why Do We Care?," NBER Chapters, in: New Developments in Productivity Analysis, pages 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, 2003*.

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*, Spring. pp.109-157.

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

Cette, Gilbert & Fernald, John & Mojon, Benoît, 2016. "The pre-Great Recession slowdown in productivity," *European Economic Review*, Elsevier, vol. 88(C), pages 3-20.

(CBO) Congressional Budget Office 2007.

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

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

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 at 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. Oxford University Press.

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 (2016a): 203–7.

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

Elsby, Michael W.L., Bart Hobijn, and Ay¸segu¨l S¸ahin, "On the Importance of the Partici-pation Margin for Labor Market Fluctuations," Journal of Monetary Economics, 2015, 72, 64 – 82.

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

Fatas, Antonio (2000). "Do business cycles cast long shadows? Short-run persistence and economic growth." Journal of Economic Growth, Vol 5, issue 2, 147-162.

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

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.

______. 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,* October. 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 Tabarrok, Alexander T. 2014. Is Regulation to Blame for the Decline in American Entrepreneurship? (December 1, 2014). GMU Working Paper in Economics No. 15-11.

Gordon, Robert J., “The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections,” Working Paper 19895, National Bureau of Economic Research, February 2014.

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

Hall, Robert and Nicolas Petrosky-Nadeau, “Changes in Labor Participation and Household Income,” Economic Letter, Federal Reserve Bank of San Francisco, May 2016, (2).

Hall, Robert E., “Macroeconomics of Persistent Slumps,” Working Paper 22230, National Bureau of Economic Research, May 2016.

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, March 2002, 92 (1), 220–239.

Jorgenson, Dale W. and Zvi Griliches (1967). “The Explanation of Productivity Change.” The Review of Economic Studies, 34(3), pp. 249-283.

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

Liu, Zheng, and Pengfei Wang. 2014. "Credit Constraints and Self-Fulfilling Business Cycles." American Economic Journal: Macroeconomics, 6(1): 32-69Müller, Ulrich and Mark Watson (2015). “Measuring Uncertainty about Long-Run Predictions.” Forthcoming Review of Economic Studies.

Müller and Watson (2015)

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

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

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) May 2013.

Polivka, Anne E. and Stephen M. Miller, "The CPS After the Redesign: Refocusing the Economic Lens," March 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, pp. 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 (2016). "Factor Models and Structural Vector Autoregressions in Macroeconomics." Forthcoming, *Handbook of Macroeconomics*, v. 2. Elsevier.

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

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, 2016.

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

Figure 1: Unemployment Rate

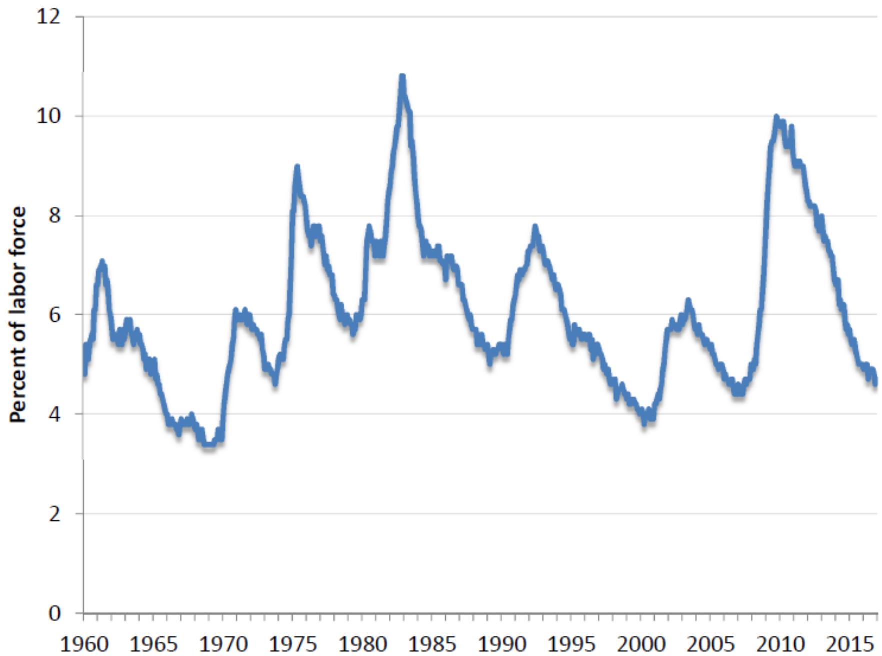

Figure 2. Survey of Professional Forecasters Median Forecasts of GDP and the Unemployment Rate, made in 2010 through 2015

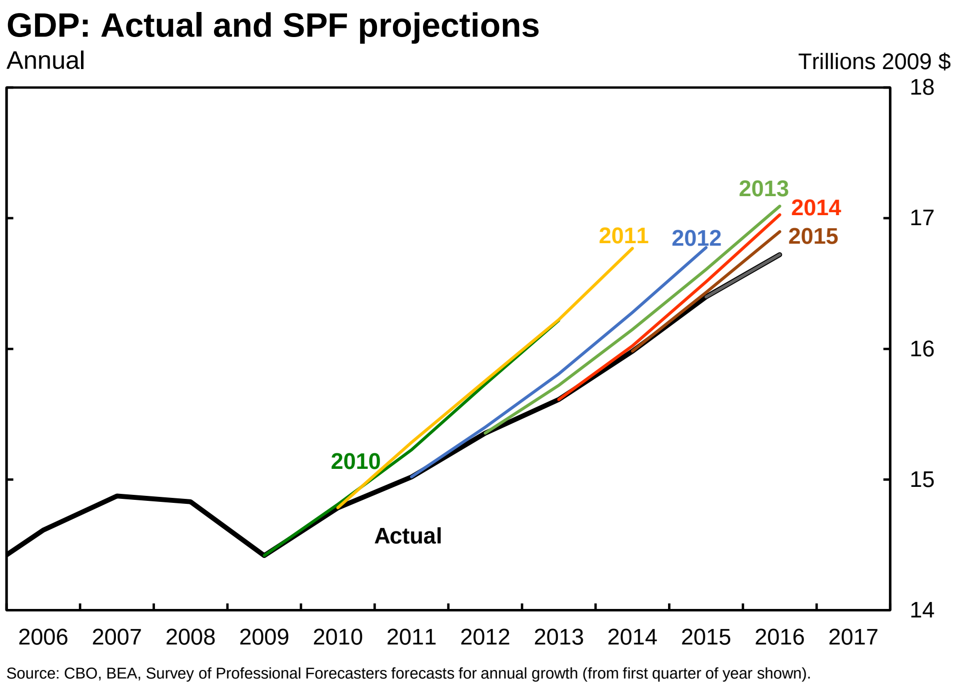

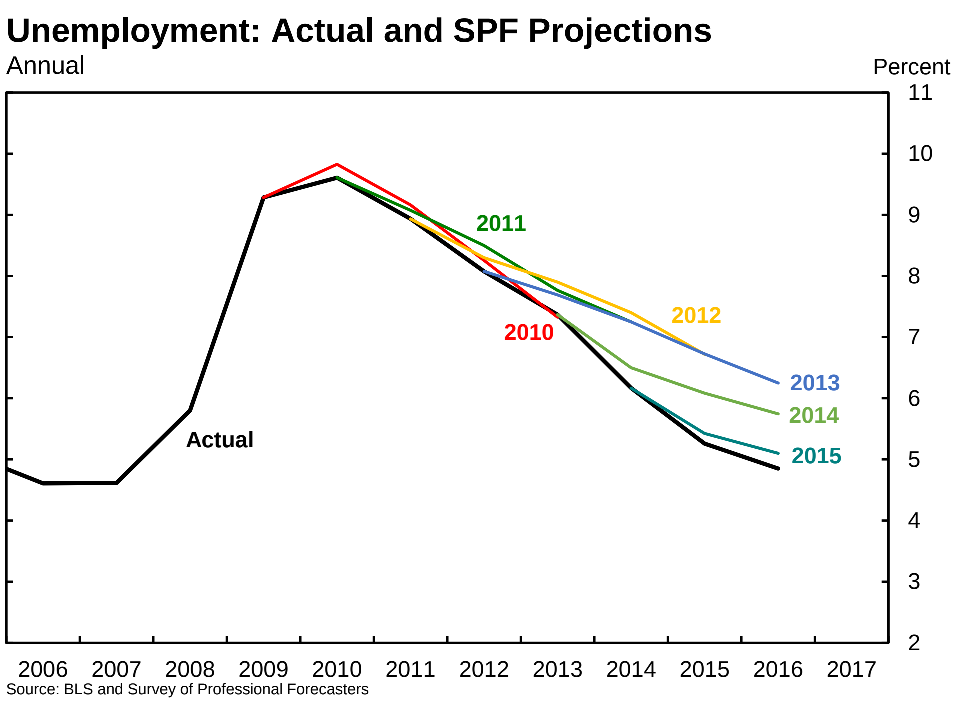

Notes: Top panel uses SPF forecasts of annual growth, and assumes the previous year's level is known. For example, the line for 2010 starts at 2009 actual, and uses forecasted annual growth for years 2010 on.

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

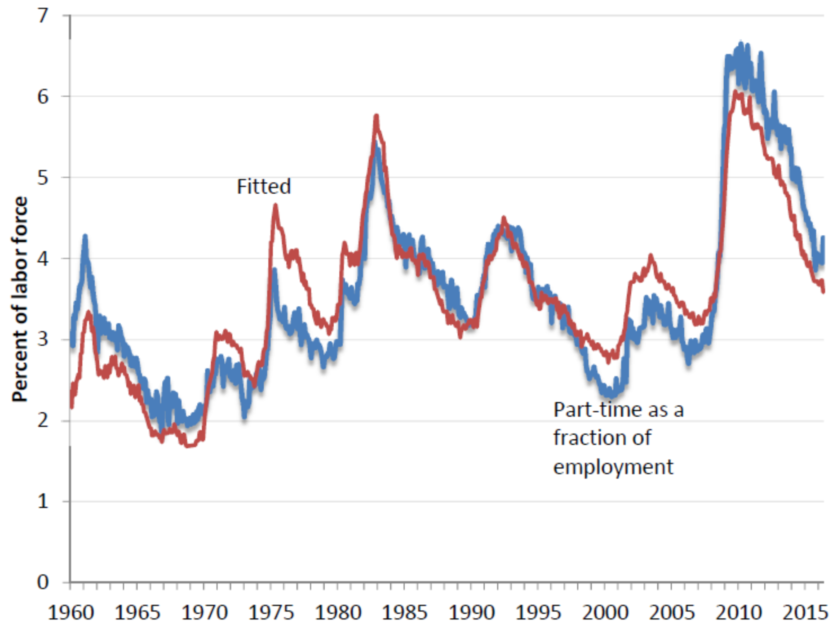

Figure 4. Growth of Low-Frequency Part of Real Business Output per Capita

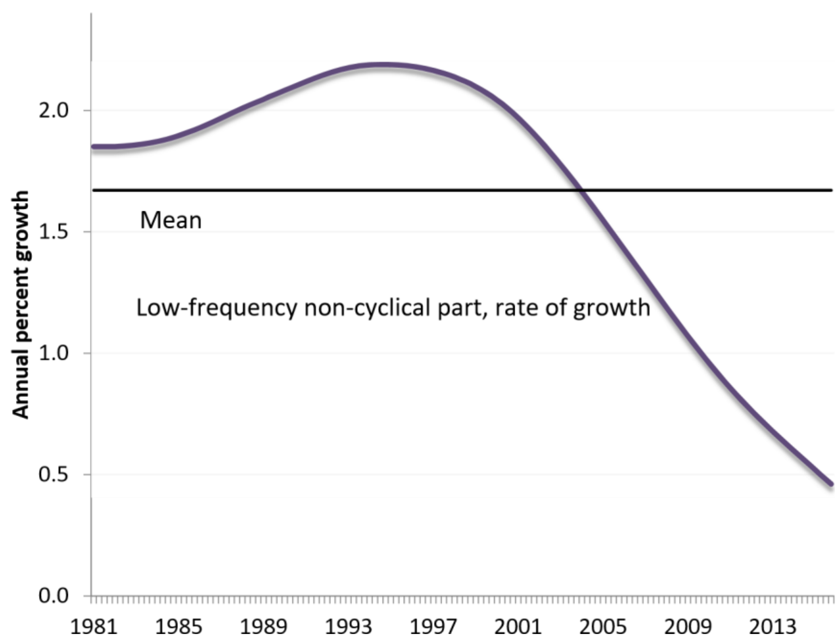

Figure 5. Levels of parts of real business output per capita

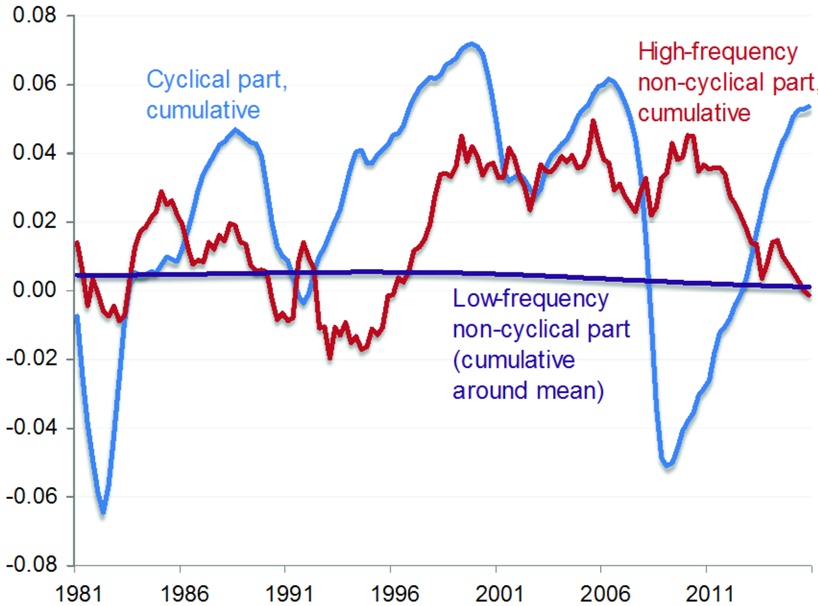

Figure 6 Growth of Low-Frequency Part of Total Factor Productivity

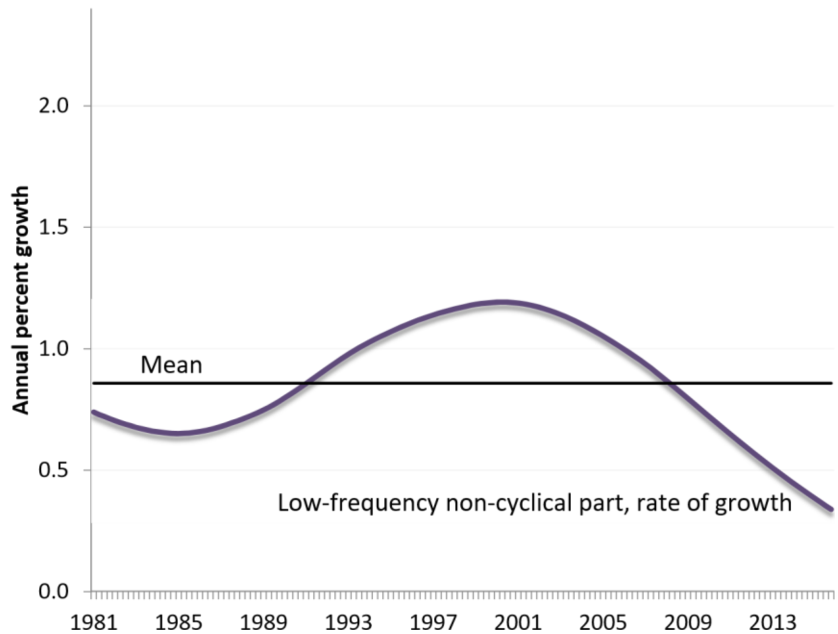

Figure 7. Levels of Parts of Total Factor Productivity

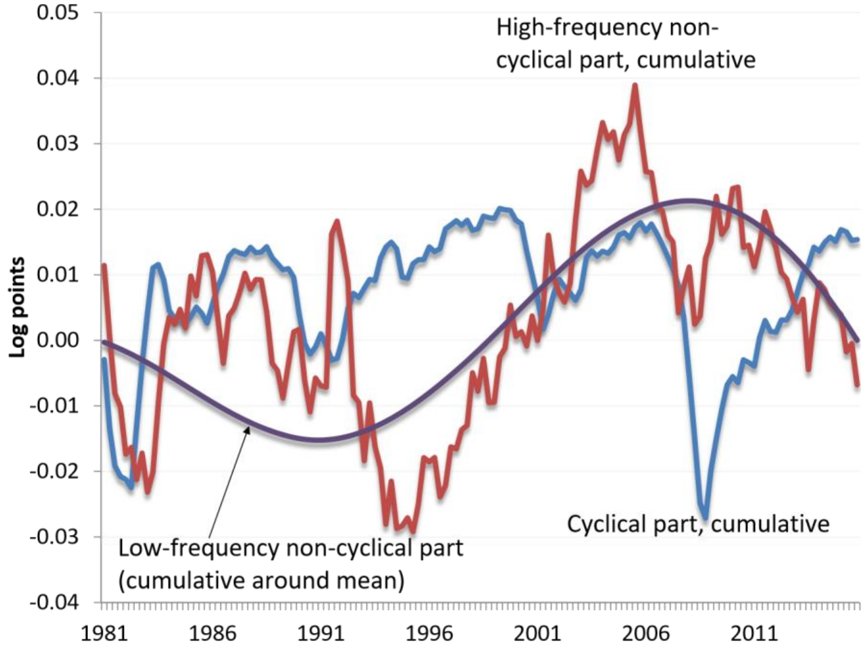

Figure 8. Levels of Parts of Weighted Capital Input

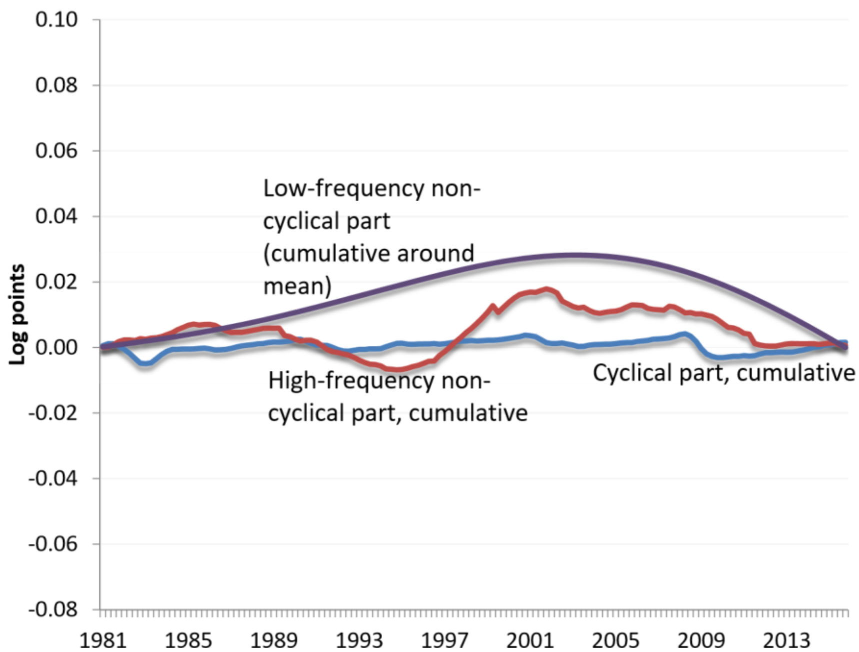

Figure 9 Levels of Parts of Hours per Worker

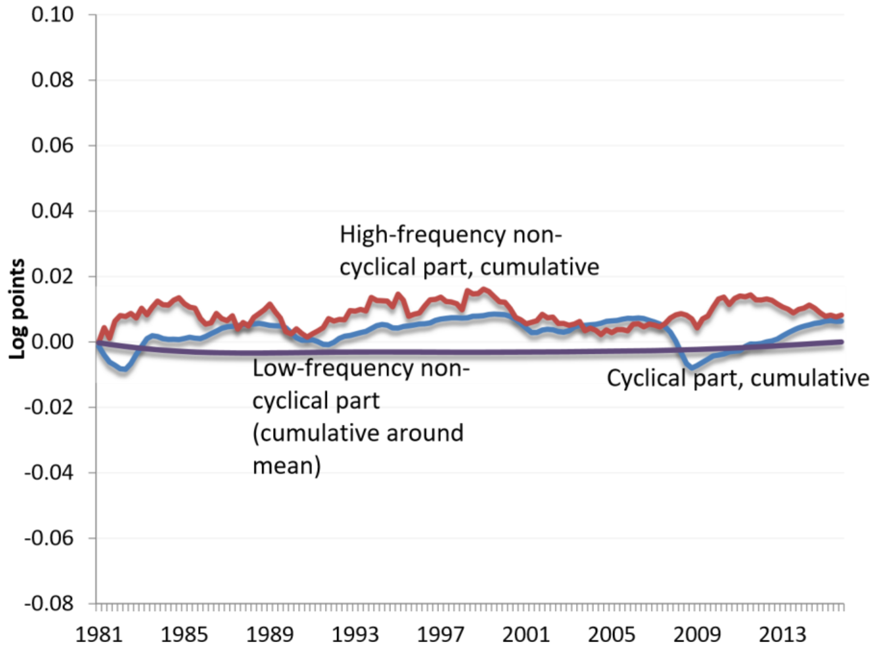

Figure 10. Growth of Low-Frequency Part of Labor-Force Participation Rate

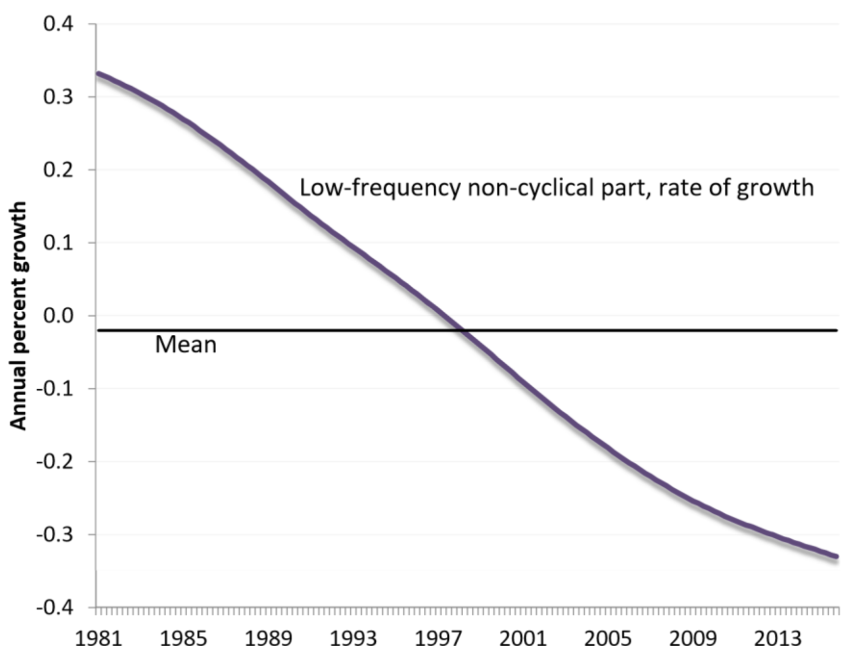

Figure 11. Levels of Parts of Labor-Force Participation Rate

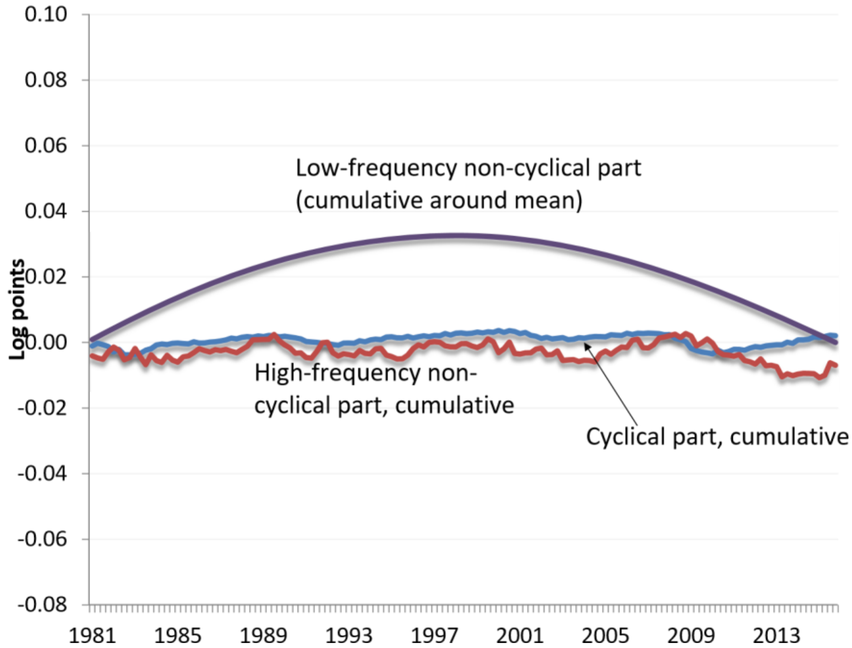

Figure 12 Levels of Parts of Labor Quality

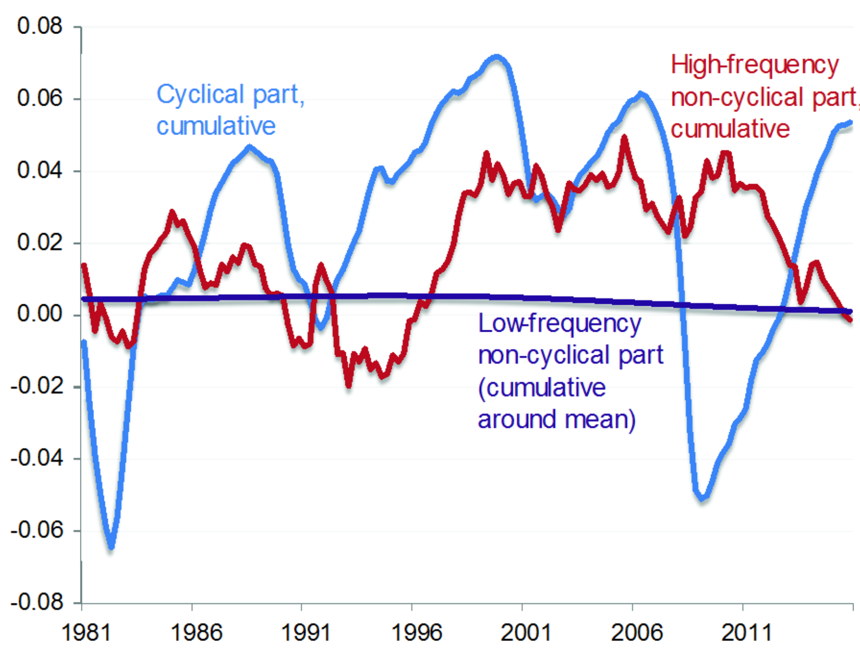

Figure 13: Historical values, trends (smooth lines), and predicted values (dashed): Growth decomposition (not share-weighted)

(a) GDP

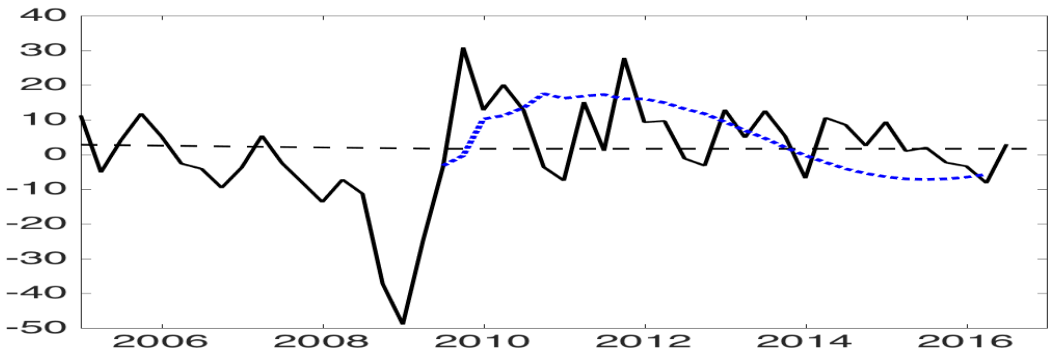

(b) GDO

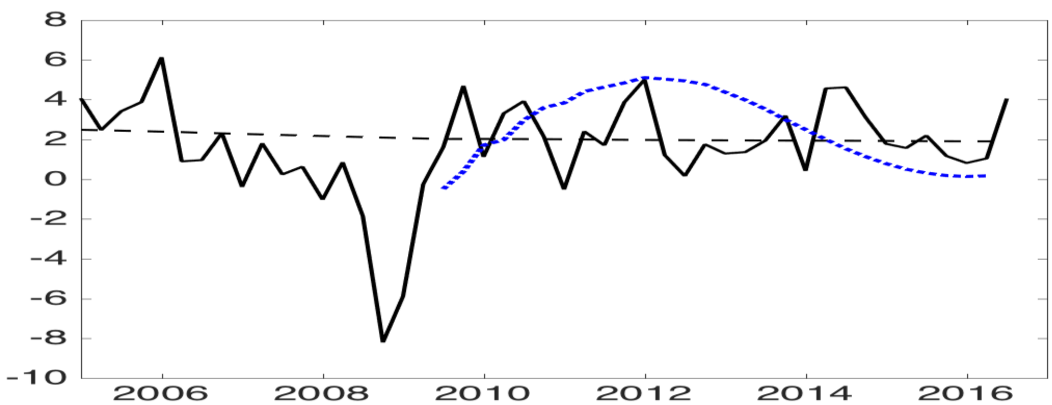

(c) GDO Business

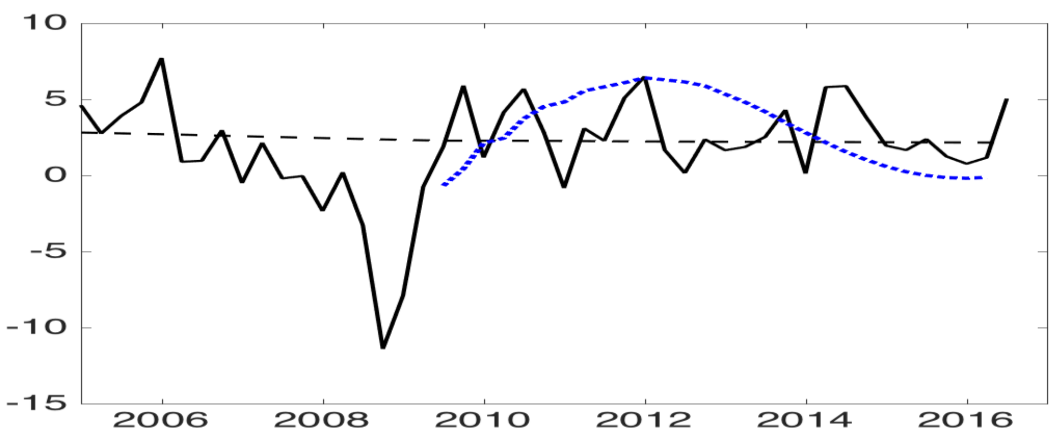

(d) productivity

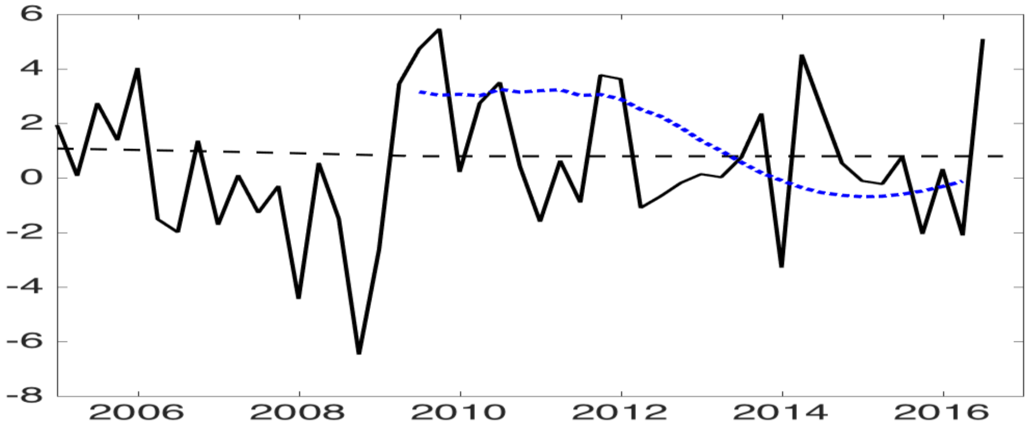

(e) Capital input

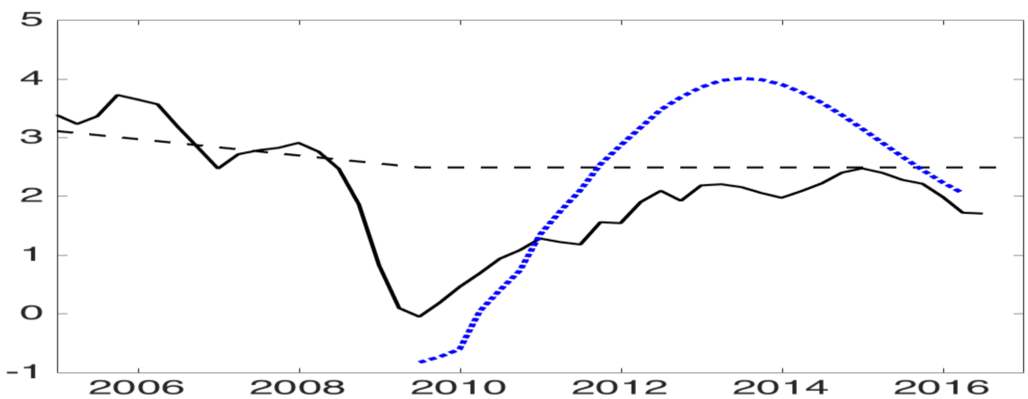

(f) Labor quality

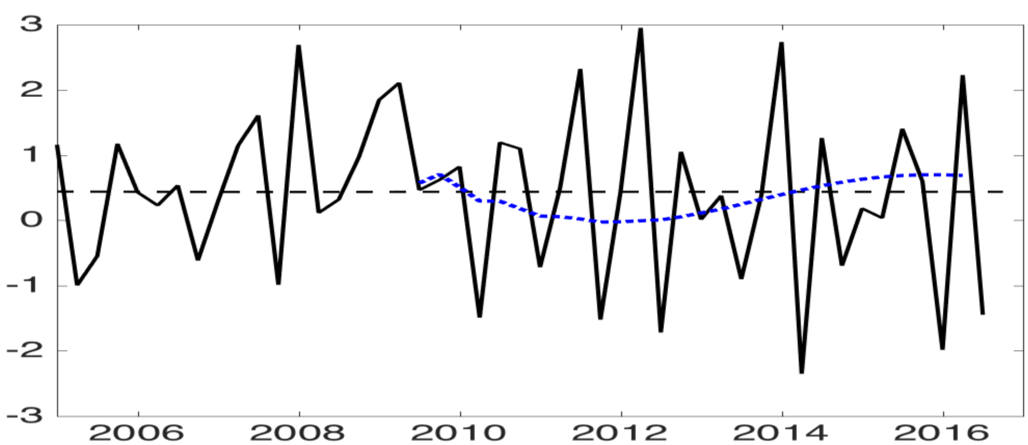

(g) Hours per worker, business

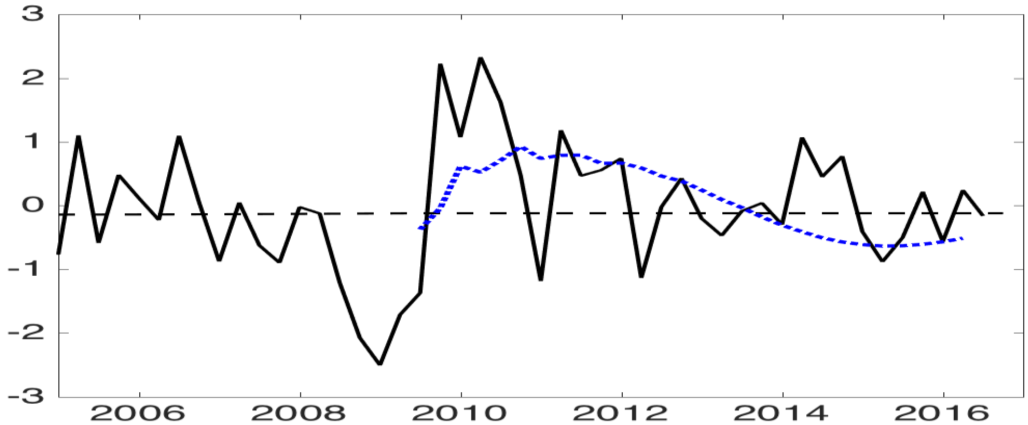

(h) Ratio, business to CPS employment

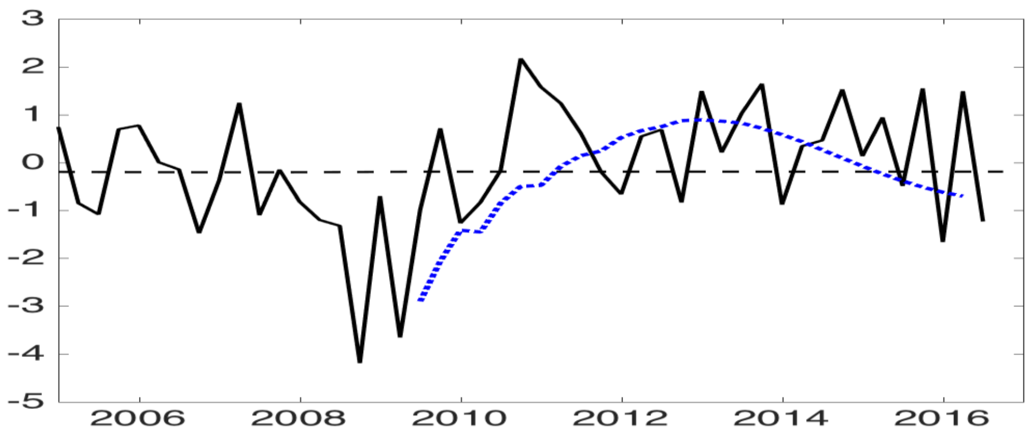

(i) CPS employment rate

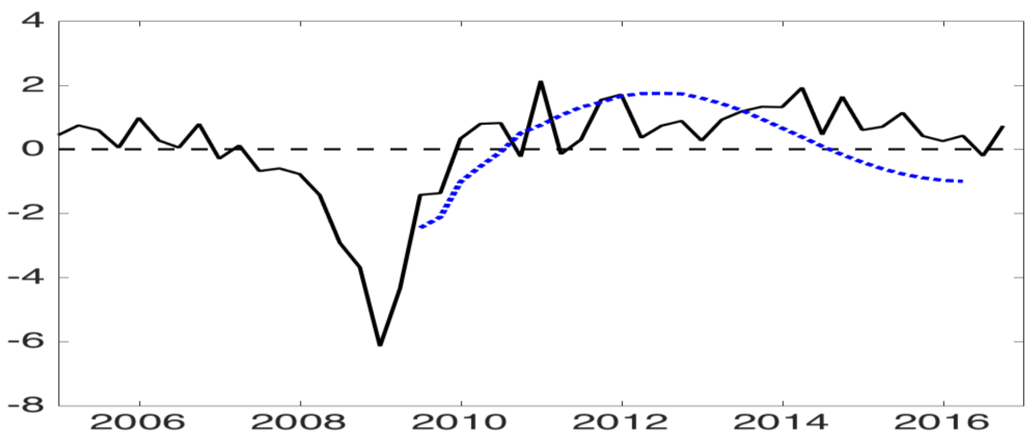

(j) Labor force participation rate

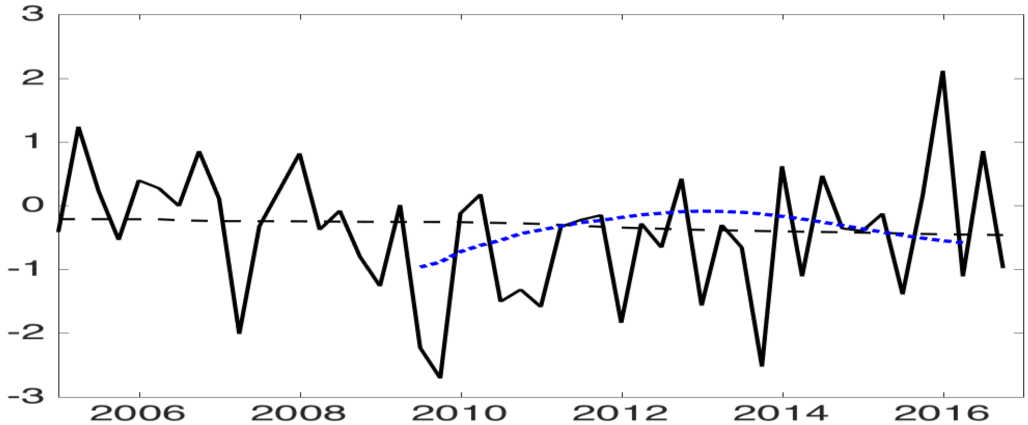

Figure 14. Investment in Productivity Improvements

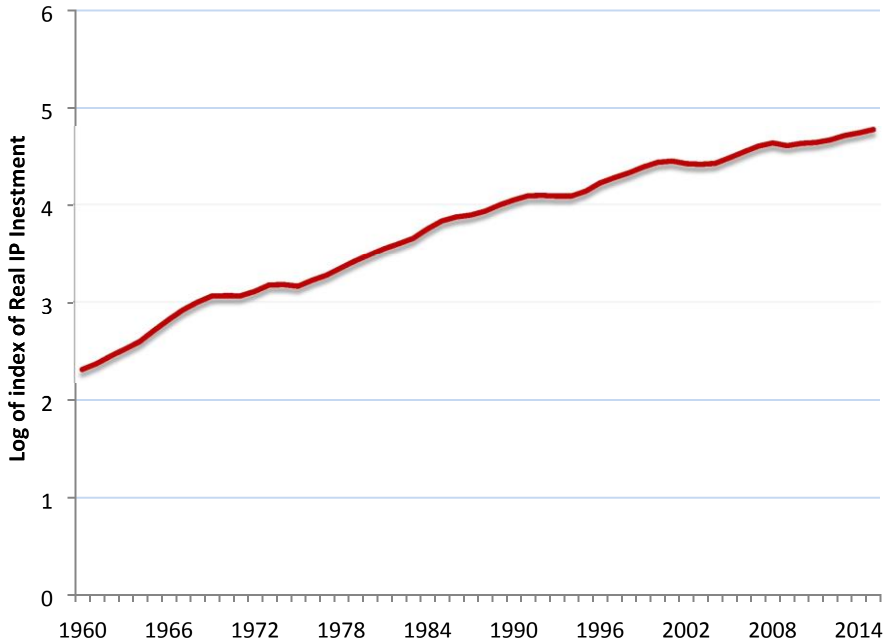

Figure 15. Equipment Investment

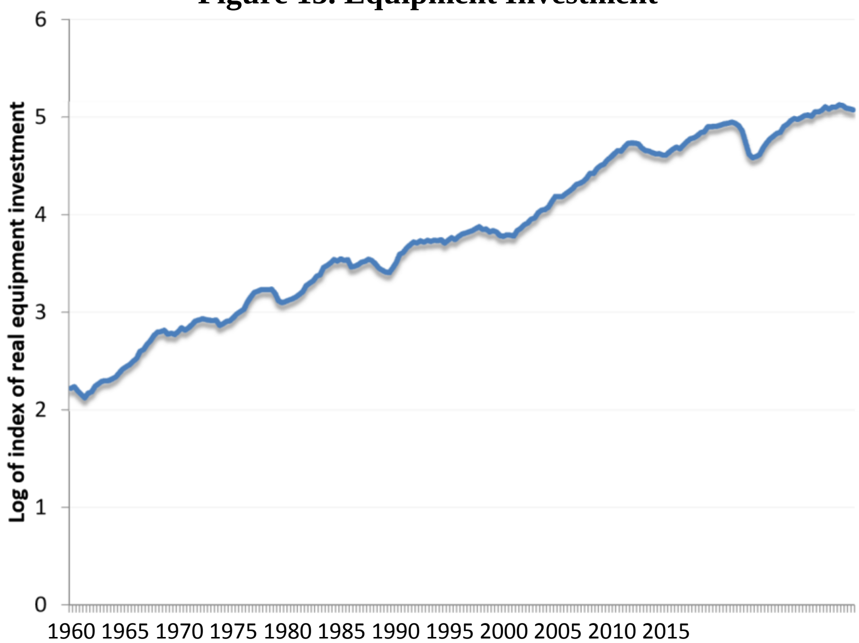

Figure 16. Business Earnings as a Ratio to the Value of Capital

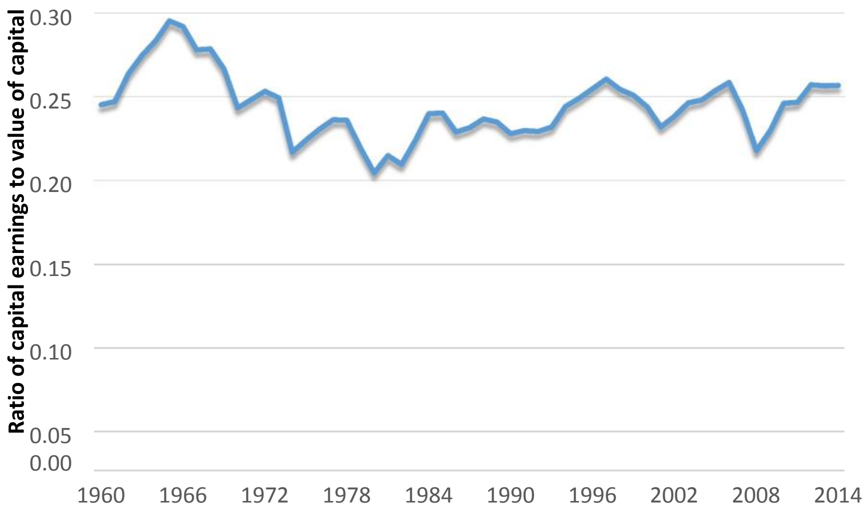

Figure 17: Real-time estimates of prospective 10-year growth in labor productivity

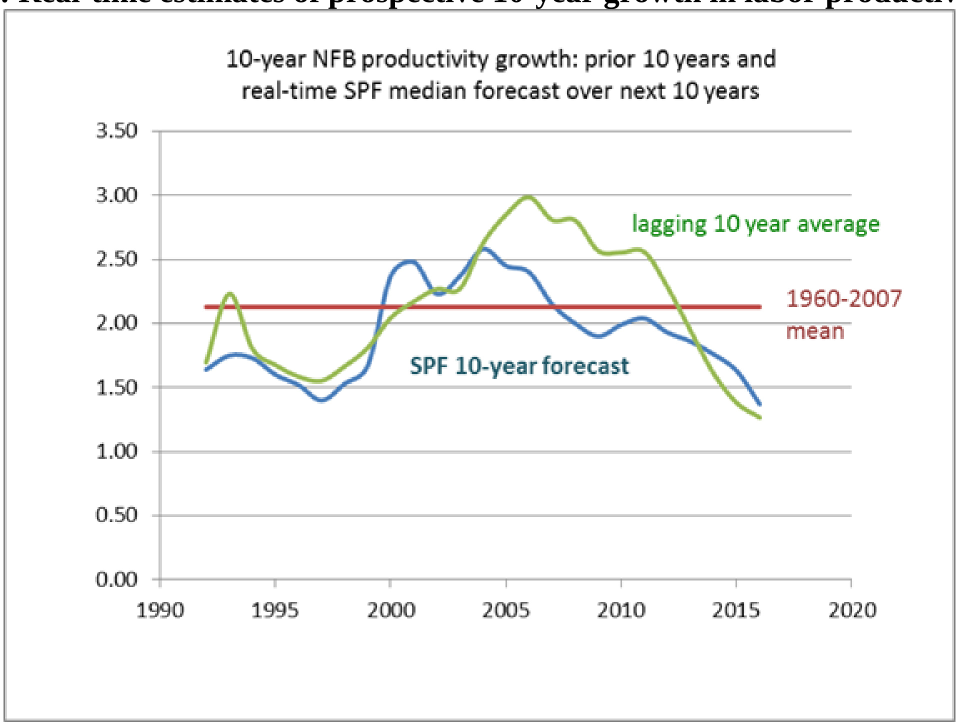

Figure 18: Levels of productivity and unemployment-adjusted productivity

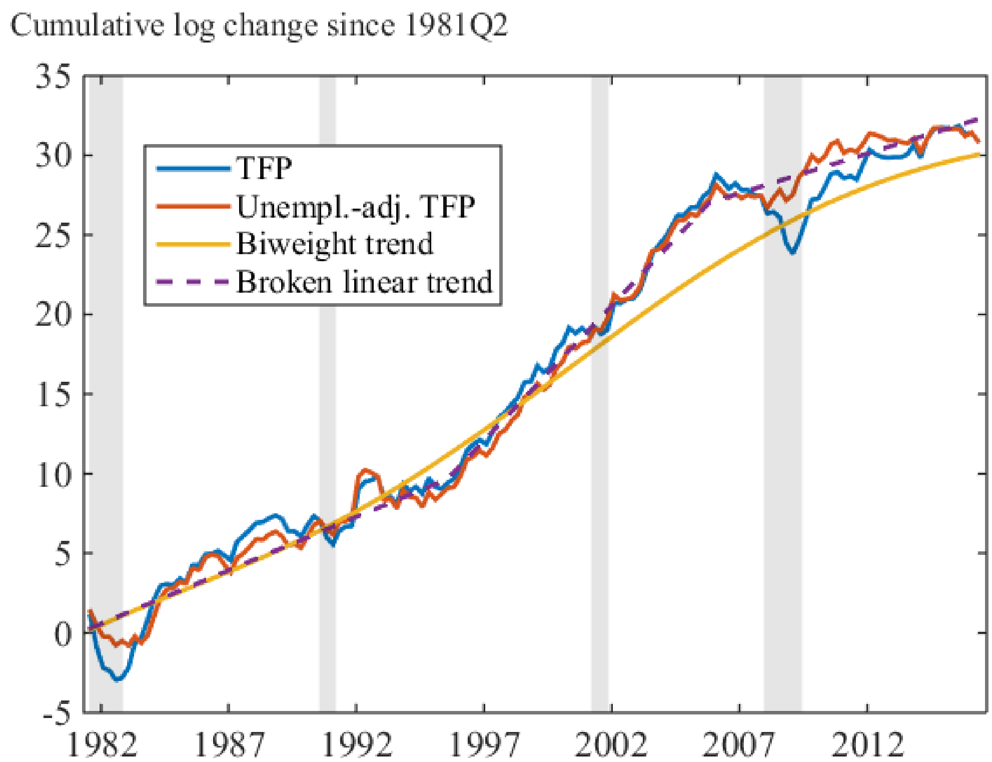

Figure 19 Labor-Force Participation Rates by Sex

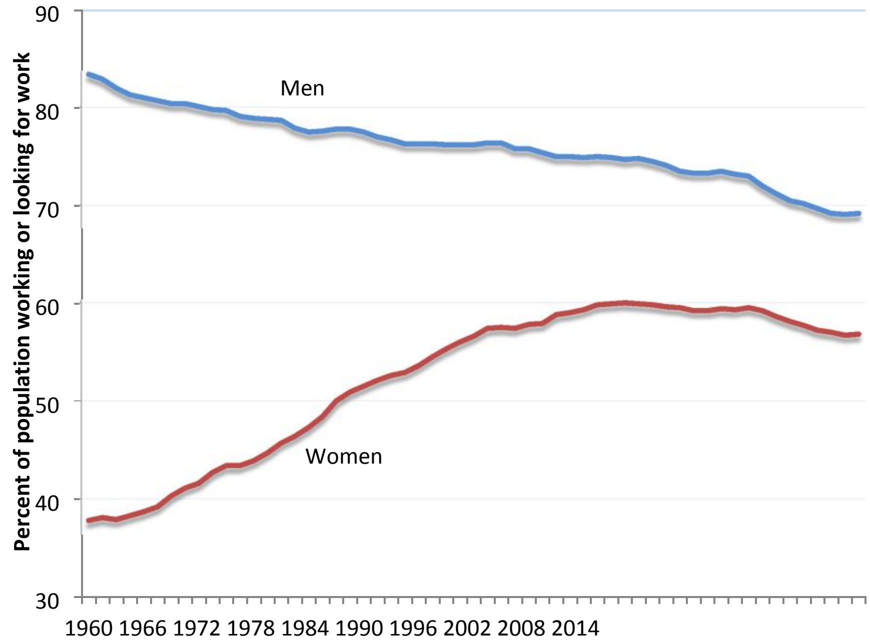

Figure 20 Effect of Education on Aging Trends in the Participation Rate since 2005

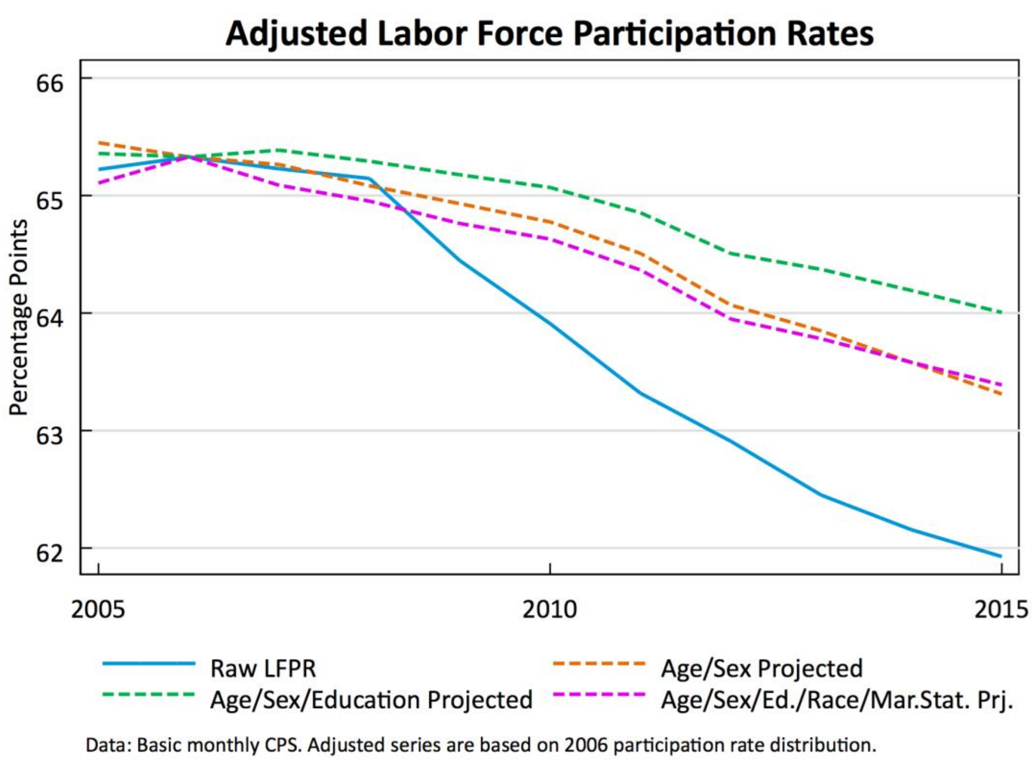

Figure 21. Role of Family Income in Participation Rates

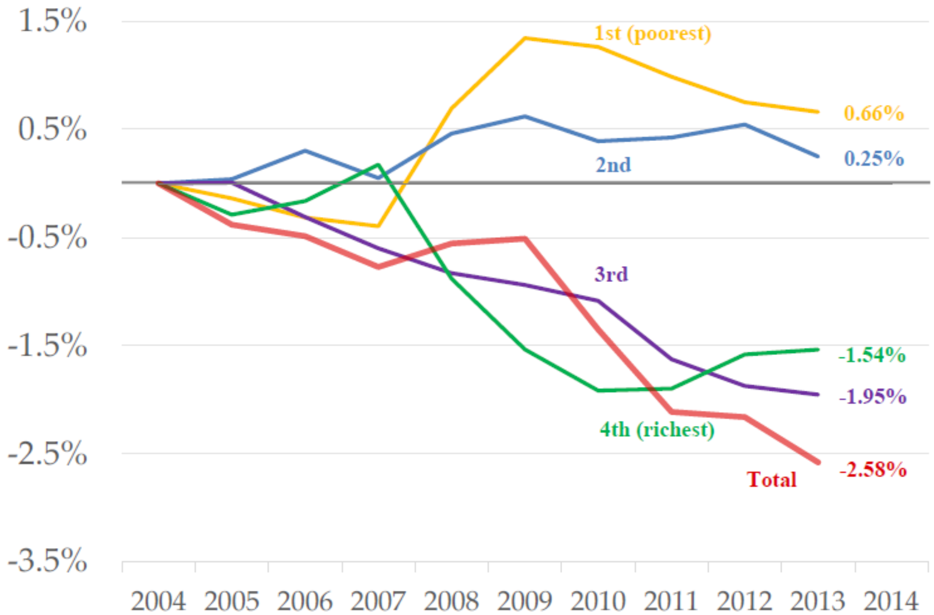

Figure 22. Historical values, trends (smooth lines), and predicted values (dashed): NIPA components (not share-weighted)

(a) Personal consumption expenditures

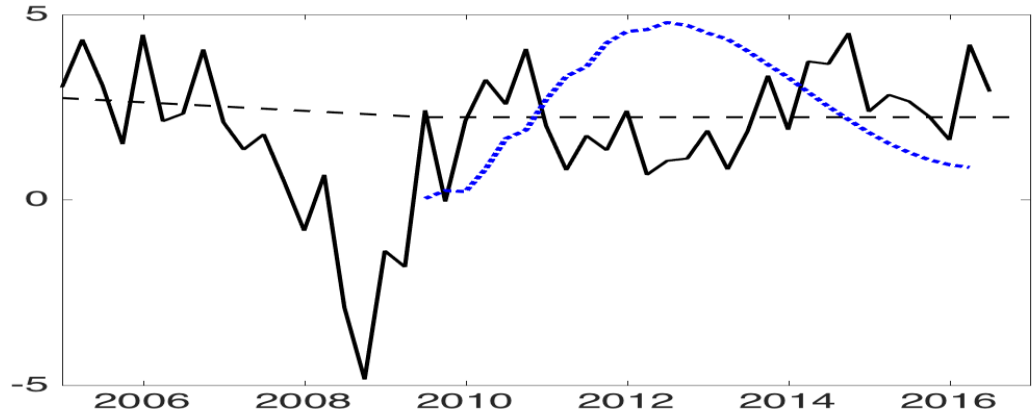

(b) PCE – durable goods

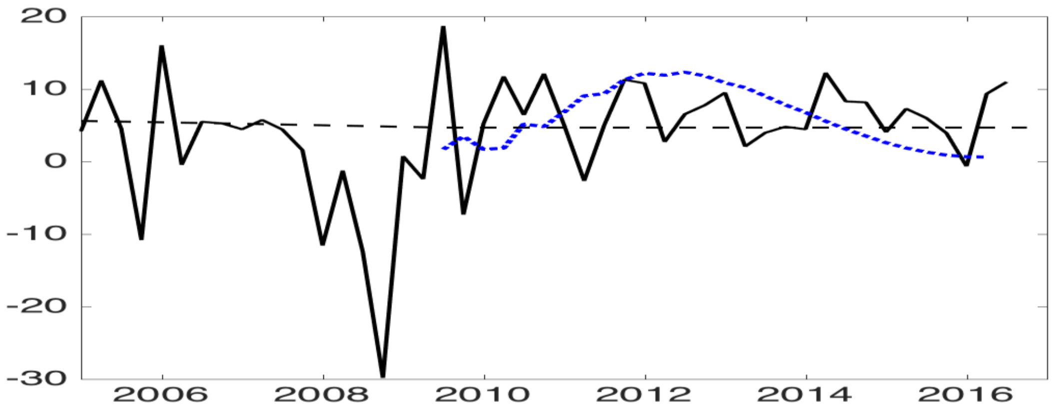

(c) PCE – nondurable goods

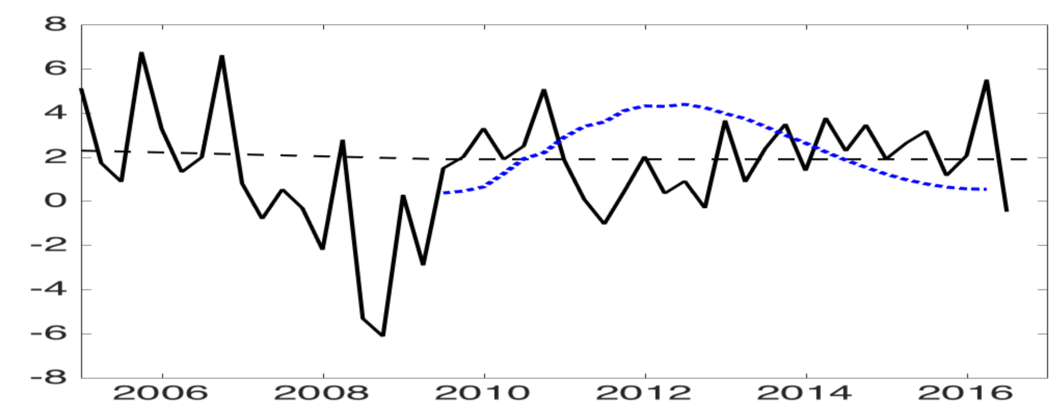

(d) PCE – services

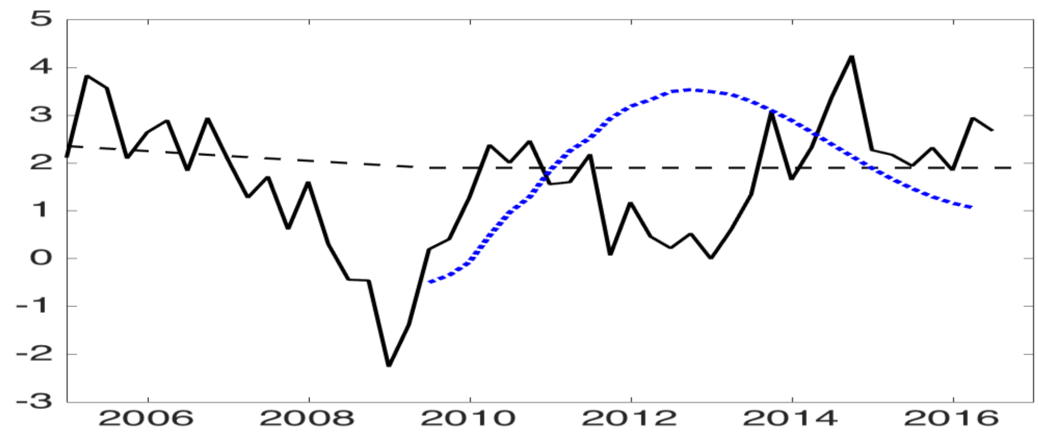

(e) PCE – services – housing & utilities

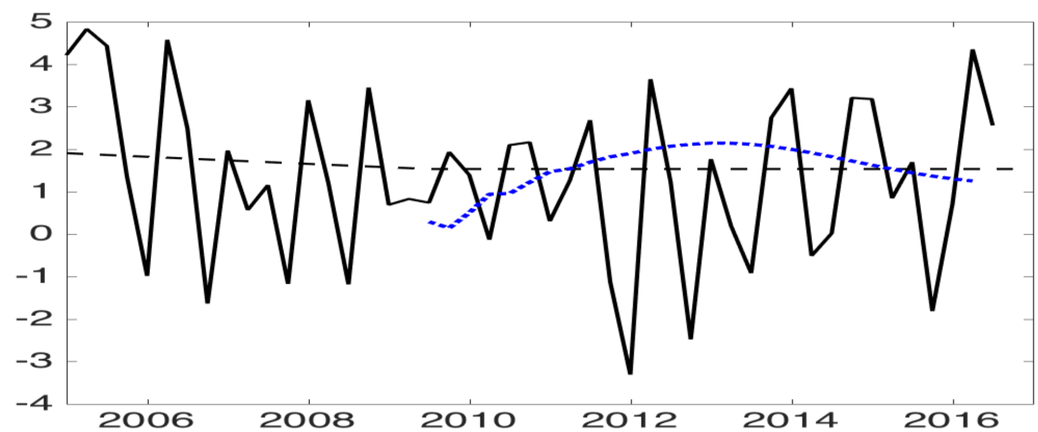

(f) PCE – services – financial services & insurance

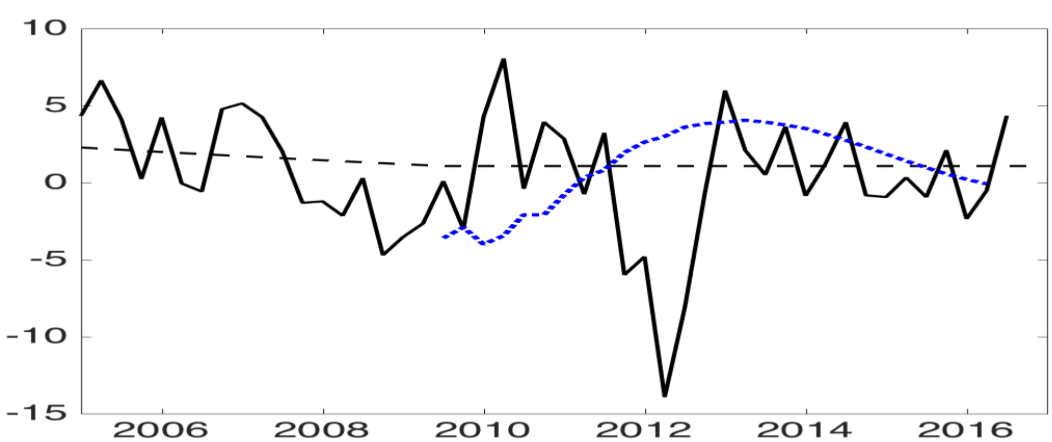

(g) Gross private domestic investment

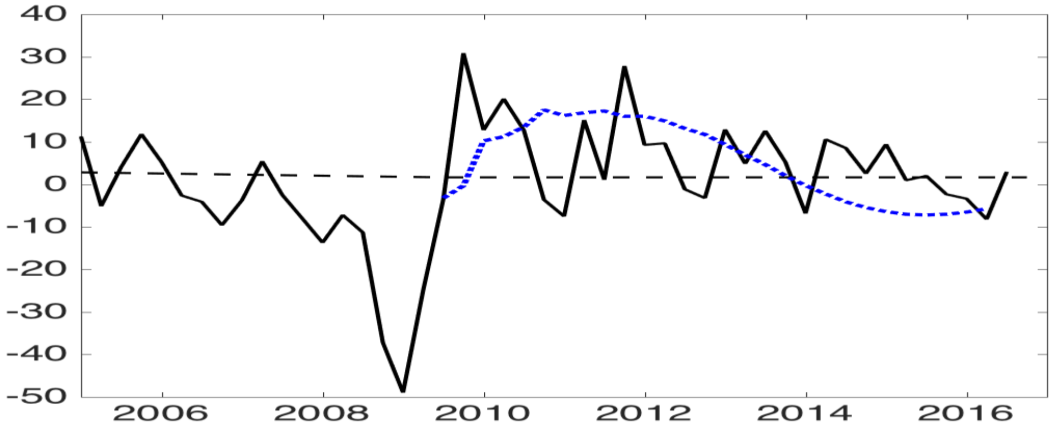

(h) Fixed private investment – nonresidential

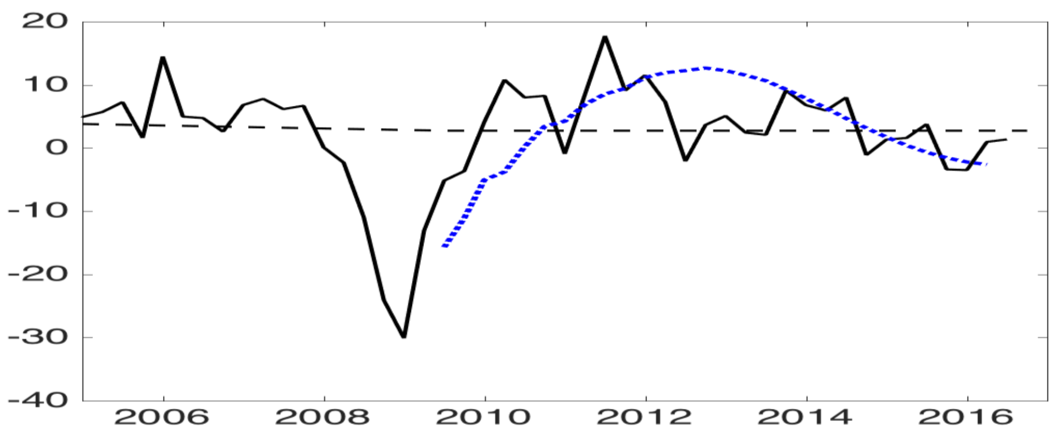

(i) Fixed private investment – residential

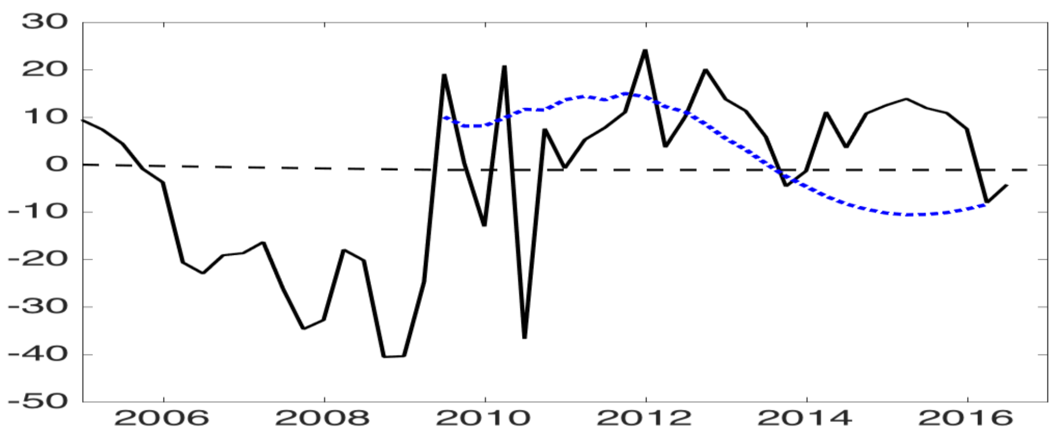

(j) Government expenditures

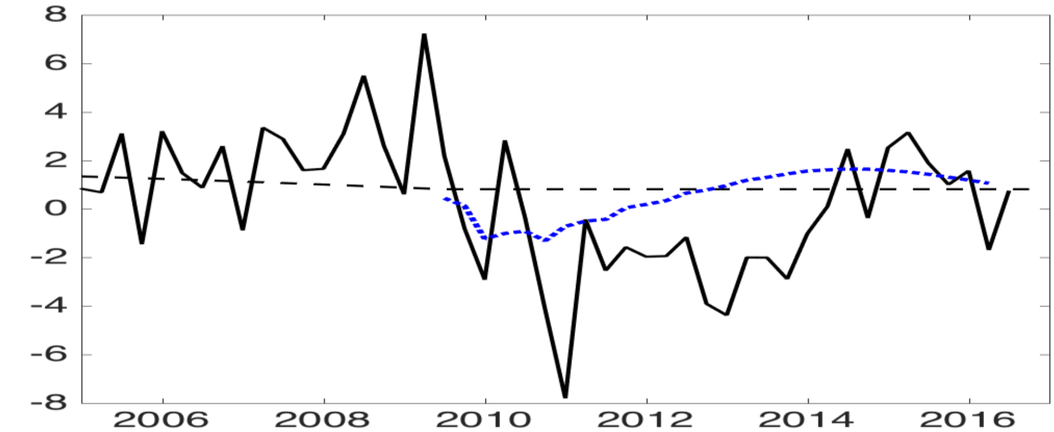

(k) Government expenditures – federal

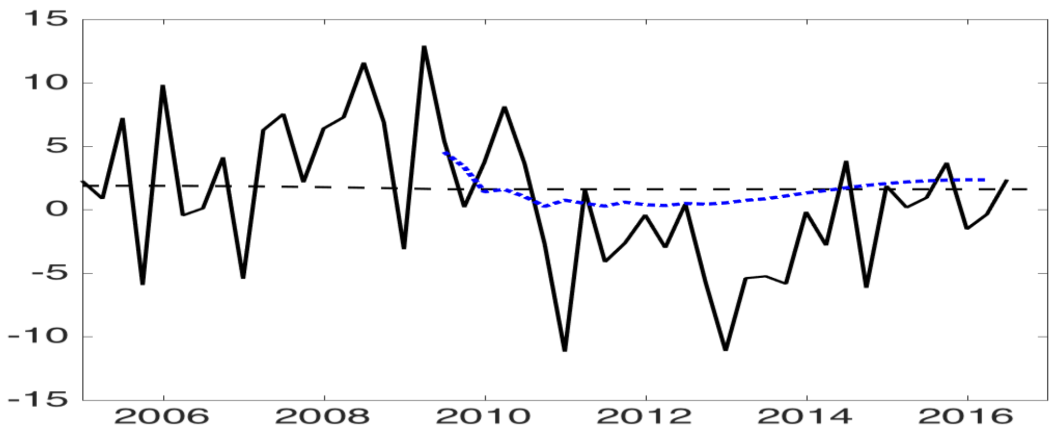

(l) Government expenditures – State & local

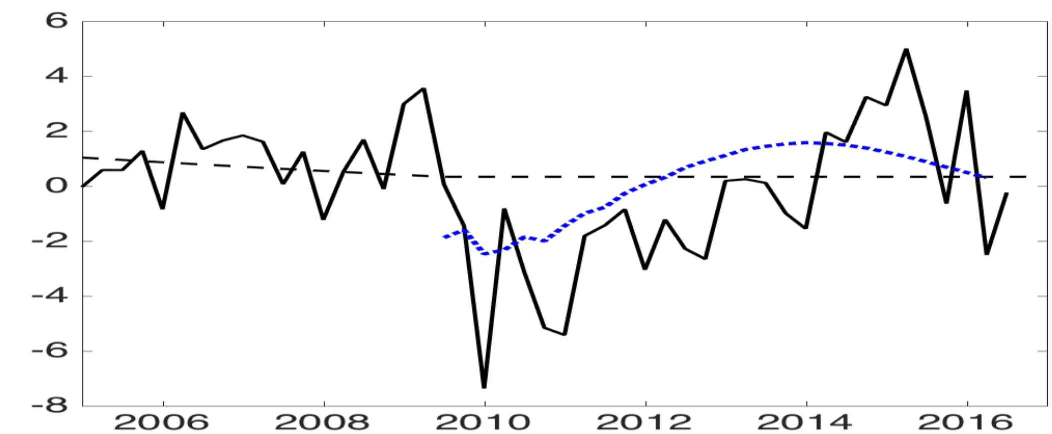

(m)Exports

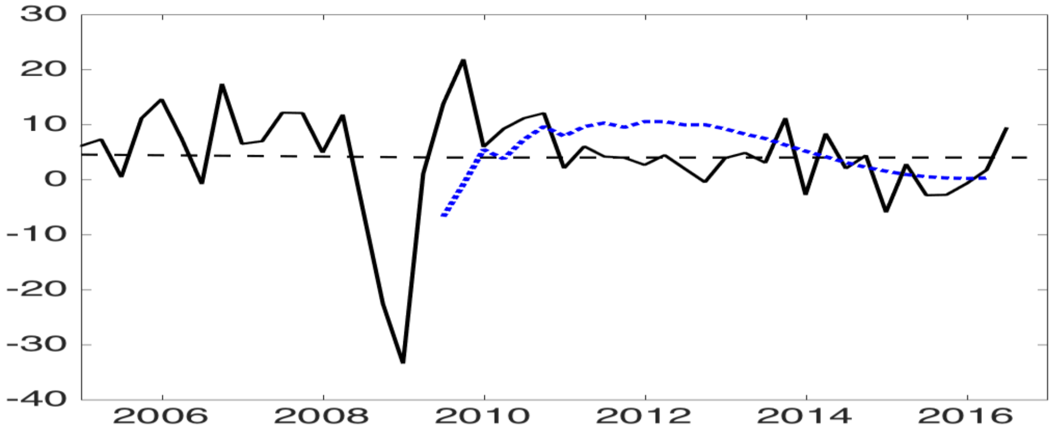

(n) Imports

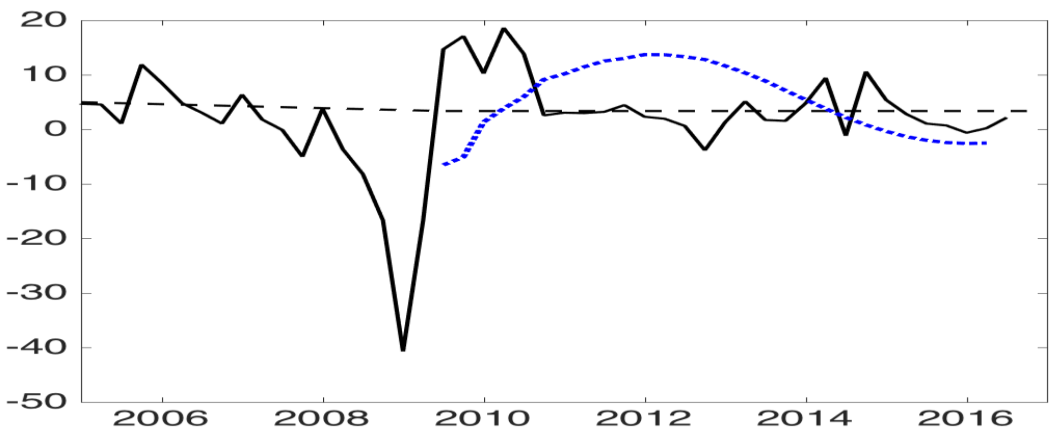

Figure 23. Historical values, trends (smooth lines), and predicted values (dashed): Additional series

(a) Employment - total

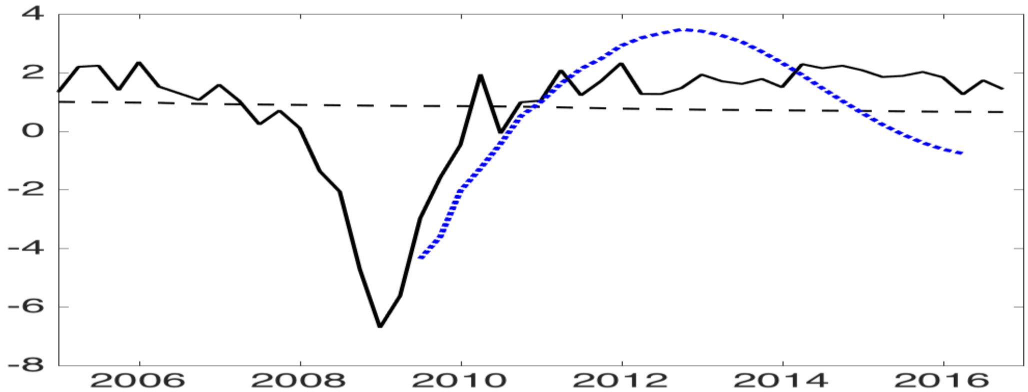

(b) Employment – business

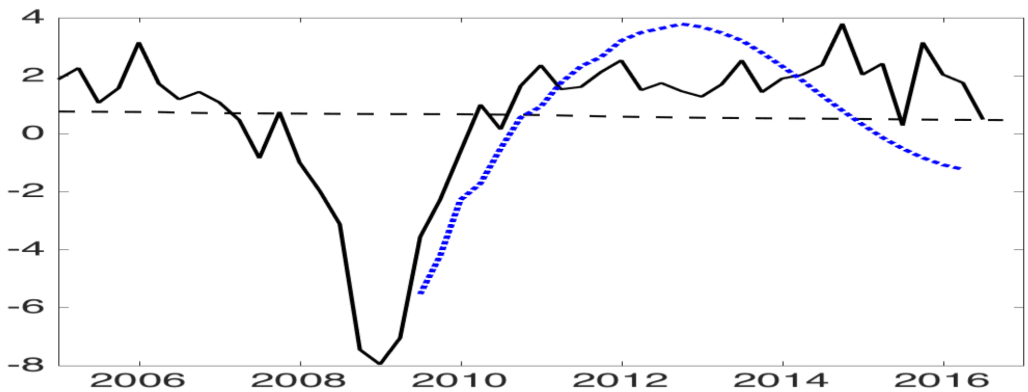

(c) Employment – state & local government

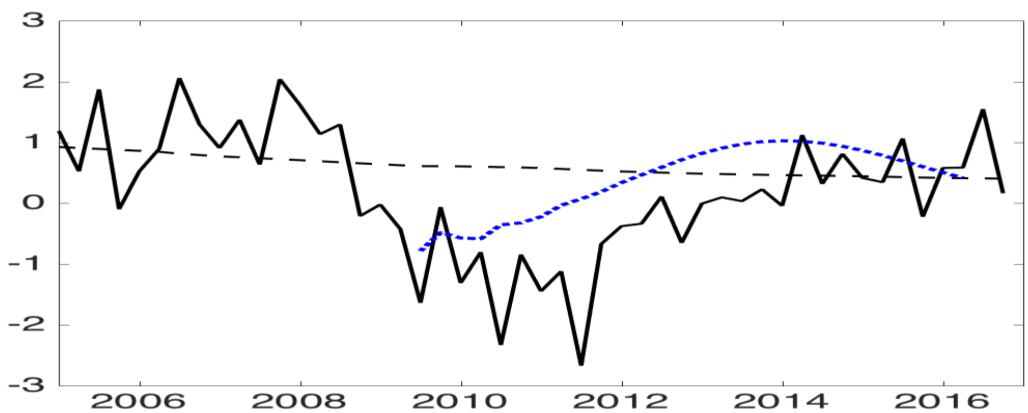

Table 1. Categories of quarterly time series used to estimate the factors

	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 (ResTable2)

Line	*Component*	*Okun's Law Estimates*						R^2 from regression of y_t onto F_t
		Generalized Okun's Law Coefficient $\beta(1)/4$ (*SE*)		*Standard Deviation of components*				
				cycle (c)	*trend* (μ)	*irregular* (z)		
1	Real gross domestic product	-1.49 (0.18)		0.71	0.19	1.59		0.66
2	Real gross domestic output	-1.53 (0.17)		0.77	0.18	1.43		0.72
3	Real business output	-2.03 (0.21)		0.98	0.21	1.89		0.73
4	Real gross domestic product per capita	-1.48 (0.17)		0.68	0.17	1.57		0.60
5	Real gross domestic output per capita	-1.52 (0.17)		0.73	0.16	1.29		0.67
6	Real business gross domestic output per capita	-2.02 (0.20)		0.94	0.18	1.73		0.70
7	Total factor productivity	-0.50 (0.19)		0.87	0.14	1.84		0.38
8	Capital input per capita × capital share	-0.09 (0.16)		0.24	0.05	0.26		0.37
9	Labor quality × labor share	0.09 (0.03)		0.20	0.00	0.78		0.06
10	Labor hours per capita × labor share	-1.52 (0.13)		1.00	0.00	0.73		0.73
11	Hours per worker, business × labor share	-0.23 (0.07)		0.15	0.01	0.56		0.25
12	Ratio of business employment to CPS employment × labor share	-0.47 (0.06)		0.43	0.02	0.65		0.24
13	CPS employment rate × labor share	-0.71 (0.01)		0.54	0.00	0.06		0.89
14	Labor-force participation rate × labor share	-0.11 (0.06)		0.26	0.02	0.58		0.02

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.

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*			
	Three previous recoveries	*Post-crisis recovery*	*Annual shortfall*	*Three previous recoveries*	*Post-crisis recovery*	*Annual shortfall*	*Cumulative percent*
Real gross domestic product	3.60	2.06	1.54	2.95	0.96	1.99	14.94
Real gross domestic output	3.57	2.20	1.37	2.92	1.11	1.81	13.54
Real business output	4.04	2.76	1.29	3.18	1.29	1.89	14.14
Real gross domestic product per capita	2.48	1.02	1.45	1.84	-0.07	1.91	14.30
Real gross domestic output per capita	2.45	1.16	1.29	1.80	0.07	1.73	12.90
Real business gross domestic output per capita	2.92	1.72	1.21	2.07	0.26	1.81	13.49
Total factor productivity	1.30	0.89	0.42	0.99	0.28	0.71	5.12
Capital input per capita × capital share	0.79	0.24	0.55	0.77	0.24	0.53	3.78
Labor quality × labor share	0.29	0.21	0.08	0.33	0.27	006	0.42
Labor hours per capita × labor share	0.54	0.38	0.16	-0.03	-0.54	0.51	3.60
Hours per worker, business × labor share	0.05	0.15	-0.10	-0.06	-0.06	-0.01	-0.06
Ratio of business employment to CPS employment × labor share	0.08	0.23	-0.15	-0.07	-0.01	-0.06	-0.42
CPS employment rate × labor share	0.29	0.42	-0.13	0.00	-0.04	0.04	0.27
Labor-force participation rate × labor share	0.13	-0.41	0.54	0.10	-0.43	0.54	3.82

Notes: Indented entries sum to the final entry at the previous indent level.

Table 4. Shortfall of the Post-Crisis Recovery Relative to 2009q2 Forecasts: Growth accounting decomposition using forecast-based cyclical adjustment

	forecast	*actual*	*forecast - actual*	*SE*
Real gross domestic product	2.57	2.06	0.51	0.07
Real gross domestic output	2.57	2.20	0.37	0.07
Real business output	3.05	2.76	0.29	0.08
Real gross domestic product per capita	1.45	1.02	0.42	0.09
Real gross domestic output per capita	1.45	1.16	0.29	0.07
Real business gross domestic output per capita	1.93	1.72	0.21	0.09
Total factor productivity	1.40	0.89	0.52	0.09
Capital input per capita × capital share	0.43	0.24	0.19	0.01
Labor quality × labor share	0.22	0.21	0.01	0.03
Labor hours per capita × labor share	-0.13	0.38	-0.51	0.04
Hours per worker, business × labor share	0.05	0.15	-0.10	0.02
Ratio of business employment to CPS employment × labor share	-0.11	0.23	-0.34	0.04
CPS employment rate × labor share	0.16	0.42	-0.26	0.01
Labor-force participation rate × labor share	-0.22	-0.41	0.19	0.02

Notes: The final column showed the standard error of the forecast associated that arises from sampling error in the estimated model parameters. It does not incorporate uncertainty from shocks, etc., over the out-of-sample forecast period.

Table 5. Shortfall of the Post-Crisis Recovery Relative to Earlier Recoveries: Labor productivity decomposition using Okun's law cyclical adjustment

	Three previous recoveries	*Post-crisis recovery*	*Annual shortfall*	*Cumulative percent*
Business labor productivity	2.12	1.01	1.11	8.09
Decomposition (5.a)				
productivity × α	0.99	0.28	0.71	5.12
Labor quality	0.49	0.43	0.06	0.44
Capital-Labor ratio × capital share	0.64	0.31	0.34	2.38
Decomposition (5.b)				
productivity / $(1-\alpha)$	1.48	0.51	0.96	6.98
Labor quality	0.49	0.43	0.06	0.44
Capital-Output ratio × $\alpha/(1-\alpha)$	0.16	0.07	0.08	0.59

Table 6. Shortfall of the Post-Crisis Recovery Relative to Earlier Recoveries: Labor productivity decomposition using forecast-based cyclical adjustment

	forecast	*actual*	*forecast - actual*	*SE*
Business labor productivity	2.11	1.09	1.02	0.08
Decomposition (5.a)				
productivity × α	1.40	0.89	0.52	0.09
Labor quality	0.34	0.33	0.01	0.04
Capital-Output ratio × capital share	0.36	-0.14	0.49	0.02
Decomposition (5.b)				
productivity / $(1-\alpha)$				
Labor quality				
Capital-Output ratio × $\alpha/(1-\alpha)$				

Table 7: Table ATUStable 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

Table 8. Expected and unexpected contributions to GDP growth: NIPA demand components

	actual	*forecast*	*forecast error*	*SE*
Real gross domestic product	**2.06**	**2.56**	**-0.51**	**0.07**
Personal consumption expenditures	**1.54**	**1.73**	**-0.19**	**0.04**
Goods	0.78	0.84	-0.06	0.03
Goods, durable	0.47	0.50	-0.03	0.03
Motor vehicles & parts	0.11	0.10	0.01	0.02
Furnishings & durable HH eqpt	0.11	0.10	0.00	0.00
Recreational goods & vehicles	0.20	0.23	-0.03	0.01
Other durables	0.05	0.06	-0.01	0.00
Goods, nondurable	0.32	0.36	-0.03	0.01
Food & beverages off premises	0.06	0.07	-0.02	0.00
Clothing & footwear	0.06	0.08	-0.02	0.01
Gasoline & energy	0.00	0.00	0.00	0.01
Other nondurable goods	0.19	0.19	0.01	0.01
Services	0.76	0.89	-0.13	0.02
Housing & utilities	0.13	0.19	-0.06	0.01
Health care	0.31	0.29	0.01	0.01
Transportation services	0.04	0.03	0.01	0.00
Recreational services	0.04	0.05	-0.01	0.00
Food services & accommodations	0.11	0.11	0.01	0.00
Financial services & insurance	0.00	0.06	-0.06	0.01
Other services	0.10	0.11	-0.01	0.01
NPISH	0.03	0.06	-0.03	0.01
Gross private domestic investment	**0.91**	**0.88**	**0.03**	**0.04**
Fixed private investment	0.70	0.57	0.12	0.03
Nonresidential	0.50	0.47	0.03	0.02
Residential	0.20	0.11	0.09	0.02
Government expenditures	**-0.19**	**0.10**	**-0.29**	**0.03**
Federal	-0.09	0.10	-0.19	0.02
State & local	-0.10	0.01	-0.11	0.01
Exports	**0.58**	**0.59**	**-0.02**	**0.04**
Imports	**-0.76**	**-0.75**	**-0.01**	**0.03**

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