

Why Has the Unemployment Rate Fared Better than GDP Growth? *

Robert E. Hall
Hoover Institution and Department of Economics,
Stanford University
National Bureau of Economic Research
rehall@stanford.edu; stanford.edu/~rehall

September 27, 2016

Abstract

Why Has the Unemployment Rate Fared Better than GDP Growth?

Answer: Between 2007 and 2014, GDP growth was held back by shortfalls of

- 4.4 percent in productivity
- 4.0 percent in capital input
- 3.6 percent in labor-force participation
- 2.2 percent in growth of the working-age population

JEL E22, E24, E32

*Prepared for the Boston Federal Reserve's conference, "The Elusive 'Great Recovery': Causes and Implications for Future Business Cycle Dynamics", October 14 and 15, 2016. This research was supported by the Hoover Institution. It is part of the Economic Fluctuations and Growth Program of the National Bureau of Economic Research. Complete backup for all of the calculations and data sources is available from my website, stanford.edu/~rehall.

In 2016, the unemployment rate finally reached normal after rising to 10 percent soon after the financial crisis of 2008. In terms of the utilization of the labor force, the performance of the U.S. economy is now good. But real GDP—and therefore real income—is far below its normal growth path. By that standard, the performance of the economy is abysmal. The clash between the two measures of performance has no parallel in the history of the economy since 1948, when accurate measurement of unemployment began.

In this paper, I decompose the shortfall of output, relative to its expected movements given unemployment, into four major categories, each corresponding to a topic widely studied by specialists. In order of quantitative importance, these are shortfalls in productivity, capital inputs, the labor force, and the working-age population. I use the tools of growth economics to carry out the quantification.

Research in this area is deeply under the influence of the legacy of Arthur Okun. His Law summarizes the normal relation between unemployment and real GDP. U.S. experience since the crisis is a major deviation from the Law. A law is a reliable empirical relation among observed variables. I show that Okun’s Law has been generally stable, so it qualifies as a law, but there is a lot to be learned from dissecting deviations around the law.

A long tradition of macroeconomics speaks of cyclical and non-cyclical movements of aggregate variables. Many macroeconomists have viewed the distinction as one of frequency—the cycle moves at frequencies up to a few years and the non-cycle is slower-moving influences, including productivity growth, population growth and movements of the ratio of the labor force to the population. This school applies filters to data to separate a high-frequency cyclical part from a low-frequency trend part. I am skeptical of the value of that approach. I believe that Okun’s approach is superior. It associates the business cycle with unemployment, so the cyclical part is the one correlated with unemployment and the non-cyclical part is the residual uncorrelated with unemployment.

1 Unemployment

The perspective of this paper is that the unemployment rate is a reliable measure of the utilization rate of labor and therefore of the cyclical state of the economy. Figure 1 shows the rate, the fraction of the labor force who are looking actively for work but are not working. Unemployment has been measured accurately and consistently since 1948. The series has little trend. Unemployment tracks the business cycle closely. It rises sharply when recession




Figure 1: Unemployment Rate

strikes, then declines to normal more gradually during the following recovery. Previously, the largest increases occurred from closely spaced recessions in 1970 and 1973-75 and again in 1980 and 1981-82. The first of these double shocks raised unemployment by slightly more than the increase from 2007 to 2010. The second resulted in a smaller increase but a slightly higher maximum value of over 10 percent. The post-crisis unemployment burden was the worst of three major experiences since 1948, but it involved neither a higher peak unemployment rate nor a slower recovery to the normal rate.

The focus on the standard measure of unemployment competes with two alternatives: First, other measures of the state of the labor market may capture unutilized labor more comprehensively than does unemployment. Second, a more general concept, the gap between actual and potential GDP, may incorporate measures from outside the labor market to deliver a clearer message about the overall state of the economy.

1.1 Alternative measures of labor utilization

The Bureau of Labor Statistics' Current Population Survey discerns three groups who are excluded from the unemployment count but are candidates for a broader measure of underutilization. The BLS describes these groups as follows: people "marginally attached to the labor force are those who currently are neither working nor looking for work but indicate

that they want and are available for a job and have looked for work sometime in the past 12 months. Discouraged workers, a subset of the marginally attached, have given a job-market related reason for not currently looking for work. [People] employed part time for economic reasons are those who want and are available for full-time work but have had to settle for a part-time schedule.” (bls.gov/news.release/empsit.t15.htm)

The BLS publishes three extended measures of labor utilization. U-4 is standard unemployment plus discouraged workers; U-5 is U-4 plus people marginally attached to the labor force; and U-6 is U-5 plus people employed part time for economic reasons.

Discouraged workers are scarce—the average value of the ratio of that category to the labor force from the time the survey included the category in 1994 was 0.36 percent. Further, the category moved over the two cycles after 1994 in parallel, after adjusting for a slight upward trend. The category tracked the rapid rise and gradual fall of unemployment that began in 2008. There is no indication that including discouraged workers in the cyclical measure would make any difference in the analysis to be described in this paper.

Marginal workers are somewhat more numerous. Their ratio to the labor force averaged 1.2 percent since 1994. This category had a small upward trend and tracked unemployment with some sign of a lag. There is a hint that marginal workers declined somewhat more slowly than unemployment around 2014, but by 2016, the two measures coincided, as indicated by a regression of the marginal/labor-force ratio on the unemployment rate. Again, almost no change would occur in the results of this paper if marginal workers were included.

The situation is different for the category of under-utilized actual workers, those holding part-time jobs when they would rather work full time. This category could matter quite a bit. It has been measured in the CPS from the start in 1948, and has averaged 3.5 percent of the labor force since then. I adjusted the data for 1948 through 1993 downward by multiplying by 0.79 to account for the change resulting from the overhaul of the CPS that went into effect in 1994 (see Polivka and Miller (1995)). It peaked at 6.5 percent soon after the crisis in 2008 and stood at 4 percent in mid-2016.

Figure 2 shows the adjusted fraction of employed people on part time for economic reasons, together with the fitted values from a regression on unemployment and a time trend. The regression tracks the data nicely for the 1981-82 and 1990 recessions. It overstates the jump in part-time in the 2001 recession, but understates it in the crisis recession, and has remained about half a percentage point above the regression throughout the recovery.




Figure 2: Fraction of Employed People on Part Time for Economic Reasons, with Fitted Value from a Regression on Unemployment

Thus there is a moderate bulge in part time work that is abnormal, by the standard of the regression. The bulge remained with variants of the regression specification.

Further research on part-time employment may uncover the causes of the recent increase in the involuntary component and help resolve the question of its relation to the measure of labor-market tightness. It seems unlikely that simply adding this component to the unemployment rate is the appropriate adjustment, given that these people are employed. For the purposes of this paper, I use the standard unemployment rate as the measure of tightness, pending development of a better measure.

1.2 The output gap

The output gap is the difference between actual and potential GDP. It is often used as a measure of under-utilized resources—that is, potential GDP is the level of output that would be achieved under a normal level of utilization. This paper investigates the reasons why potential GDP has grown so slowly during the period starting in 2010 when unemployment began to drop and eventually reached normal levels. Many of the topics are unsettled, meaning that there is a spectrum of opinions about the course of potential GDP and the resulting output gap. It would be illogical to take any specific measure of potential GDP

as the starting point for this paper. Rather, the paper is about the large changes in the determinants of potential GDP since the crisis.

2 Evidence on the Relation between Output Growth and Unemployment Changes

Okun (1962) launched the empirical investigation of the relation between output and unemployment fluctuations. He found a regression coefficient of -0.30 with the change in unemployment as the left-hand variable and the percentage change in real GNP as the right-hand variable. He inferred that the reciprocal of this coefficient, 3.3, was the expected growth of GNP for each percentage point decline in unemployment. His regression coefficient says that the expectation of the change in unemployment conditional on GNP growth is -0.30 times GNP growth. The expectation of GNP growth conditional on the change in unemployment is the R^2 of the regression times the reciprocal of the regression coefficient, or $0.62 \times -1/0.30 = -2.1$. Okun's Law is often formulated in this way. The finding of a relation of real GNP (or now, GDP) growth of two percentage points per percentage point of unemployment has been a stable feature of U.S. data since the beginning of modern unemployment measurement in 1948. The data never supported Okun's figure of 3.3 points of GNP growth per percentage point of added unemployment, in the sense of conditional expectation.

Okun interpreted the finding of large changes in output accompanying smaller changes in unemployment in terms of indirect influences on top of the direct relation arising from employment's role as an input to the production of output. These influences included induced changes in labor-force participation and in hours per worker. He also noted changes in productivity associated with what we now call overhead labor or labor hoarding.

In the years since the financial crisis of 2008, Okun's Law has failed quite seriously. Robust declines in unemployment from the peak of 10 percent in 2010 to below 5 percent have not come with extra real GDP growth cumulating to 10 percent. If they had, real GDP would be back on its pre-crisis growth path. But real GDP has grown only at roughly normal rates and has not closed the gap below trend that opened up following the crisis.

The starting point for this paper is to develop an economic framework for studying the issues that Okun raised—the induced changes in output that arise from factors apart from increased use of labor for production when the labor-market tightens. Growth theory

provides a framework for quantifying the contributions of capital and labor. Solow (1957) observed that the log-linearized production function satisfied the following equation if firms minimized cost and markets are competitive:

$$\Delta \log \text{ private output} = \Delta \log \text{ total factor productivity} + \text{capital share} \times \Delta \log \text{ capital input} + \text{labor share} \times \Delta \log \text{ labor input}$$

The measurement of total factor productivity generally focuses on the private economy. Solow's framework does not plausibly apply to government. I make use of John Fernald's compilation of productivity data for the private economy at a quarterly frequency—see Fernald (2012) and its accompanying spreadsheet. I will concentrate on issues relating to the measurement of labor input and use Fernald's data directly for output and capital input for the private economy. Fernald's concept of labor input is private hours adjusted for labor quality:

$$\log \text{ labor input} = \log \text{ private hours} + \log \text{ labor quality}$$

I use data for the total economy to break down private hours, so I make use of the identity

$$\log \text{ private hours} = \log \frac{\text{private hours}}{\text{total hours}} + \log \text{ total hours}$$

Then to focus on the role of hours per worker, I use the identity,

$$\log \text{ total hours} = \log \frac{\text{total hours}}{\text{employment}} + \log \text{ employment}$$

The direct effect of unemployment operates through the employment rate, which is $1 - \text{the unemployment rate}$:

$$\log \text{ employment} = \log \frac{\text{employment}}{\text{labor force}} + \log \text{ labor force}$$

Labor-force participation enters via the identity

$$\log \text{ labor force} = \log \frac{\text{labor force}}{\text{population} \geq 16} + \log \text{ population} \geq 16$$

Pulling all of these together yields the decomposition of output growth,
Rate of growth of output = the sum of

- the rate of growth of total factor productivity
- the capital share \times the rate of growth of the capital stock

plus the labor share \times the sum of the rates of growth of

- the quality index of workers
- the number of people 16 and over
- the fraction of people 16 and over participating in the labor force
- the fraction of those in the labor force who are employed
- the average number of hours per worker in the total economy
- the fraction of hours in the total economy that are in the private economy

I have developed the decomposition in terms of total factor productivity, a framework that is most closely affiliated with production theory, where the index of productivity is the Hicks-neutral index of technical change. Some discussions use the alternative framework of labor productivity, measured as output per unit of labor input. One can map back and forth between the two frameworks, but that mapping is outside the scope of this paper.

One aspect of Fernald's data will be important in the later discussion: Tracking the current version of the National Income and Product Accounts, Fernald measures the capital stock to include research and development and other intangible capital, along with plant, equipment, and inventories. His measure of total factor productivity is the Solow residual, the difference between actual output growth and the weighted sum of inputs, including the services of intangible capital. If treating R&D as a kind of capital takes proper account of its contributions, the resulting measure of total factor productivity would not respond to the ups and downs of R&D investment. Any change in measured productivity associated with R&D is an indication that technical improvements from R&D change productivity itself as well as contributing to output through the capital stock. Modern growth theory often hypothesizes an external effect of improvements originating in one firm, as the ideas in the improvement enter the public domain and are adopted widely for free.

This decomposition makes it possible to investigate the issues that concerned Okun and have remained important in research on cyclical fluctuations. The basic idea is that the regression coefficient of $\Delta \log$ real GDP on Δ unemployment is the sum of the coefficients on all the components in the list above. Regression is a linear operator.

The set of regressions reveals two important properties of each component. One is the extent of the relation between unemployment and that component. For example, the coefficient of the regression of the labor-force participation rate on unemployment measures the cyclical movement of participation. The second is the magnitude of movement of the component that is *not* correlated with unemployment and is therefore not cyclical. The standard deviation of the residual from the regression measures that magnitude.

3 Results

Data for private real GDP, total factor productivity, capital input, private hours of work, and labor quality are from Fernald’s spreadsheet of May 2016. Data on population, unemployment, and employment are from the Current Population Survey. Total economy-wide hours of work are from the National Income and Product Accounts. For further details, see the supporting spreadsheet for this paper.

Table 1 shows the results of the regressions that separate total output and its components into a cyclical part and a non-cyclical part. Each row is a regression of the quarterly change of the log of a component on the change in the unemployment rate in that quarter, stated as a decimal. The first two columns show the coefficients on the change in unemployment and the standard errors of the coefficients. A negative coefficient indicates a pro-cyclical component that declines when unemployment rises.

As derived earlier, the components related to factor inputs are multiplied by the corresponding factor shares, treated as estimates of the elasticity of the production function with respect to the component. Total factor productivity comes in without any multiplier. Capital input is multiplied by the capital share and all the labor-related components are multiplied by the labor share.

3.1 Implications of the regression coefficients

The first line nicely replicates Okun’s regression. After 56 years of further experience, the conditional expectation of the increase in GDP per percentage point decrease in unemployment remains almost exactly what Okun found, a bit above two percent. As Okun recognized, this is a big number.

The remaining lines in the table allocate the coefficient of -2.1 to the contributions of the components derived earlier. Because the components multiply together to match private

<i>Line</i>	<i>Component</i>	<i>Regression coefficient on unemployment rate</i>	<i>Standard error</i>	<i>Cyclical standard deviation</i>	<i>Non- cyclical standard deviation</i>
1	Private real GDP	-2.125	(0.128)	3.33	3.26
2	Total factor productivity	-0.911	(0.124)	1.43	3.15
3	Capital input	-0.032	(0.015)	0.05	0.38
4	Population 16 and over	0.018	(0.017)	0.03	0.44
5	Labor-force participation rate	0.025	(0.033)	0.04	0.85
6	Employment rate	-0.722	(0.001)	1.13	0.03
7	Hours per worker	-0.516	(0.074)	0.81	1.89
8	Ratio of private to total hours of work	-0.051	(0.056)	0.08	1.43
9	Labor quality	0.063	(0.021)	0.10	0.52

Notes: Components are first-differences of logs. The unemployment rate is in first-differences. The cyclical standard deviation is 4 times the standard deviation in percentage points of the fitted value in the regression and the non-cyclical standard deviation is 4 times the standard deviation of the residual. Each row is a separate regression with only one right-hand variable. Data run from 1948 second quarter to 2014 fourth quarter.

Table 1: Regression Results for Real GDP and Its Components

real GDP exactly (and so the logs add up to the log of GDP) and because regression is a linear operator, the coefficients for the components add up to exactly the coefficient in the first line, -2.125 .

Okun identified two primary sources of additional cyclical behavior of output: productivity and hours. Line 2 of the table shows that total factor productivity is the single biggest indirect contributor to the high cyclical behavior of real GDP. An increase of one percentage point in unemployment goes with a decrease of -0.911 percentage points of total factor productivity. Research on productivity has discussed the roles of labor hoarding, cyclical changes in capital utilization, measurement errors, and other factors that account for the pro-cyclical behavior of productivity. Fernald has an alternative measure of TFP that separates a component called capital utilization. For the sake of simplicity, I use his measure that includes that component within total factor productivity. The results would be identical if TFP and Fernald's utilization adjustment were separated—separation would just add another component.

Line 3 shows that cyclical effects on capital input has essentially no role in cyclical movements of output. Though investment is pro-cyclical, the cumulated stock of capital changes relatively little in synchrony with unemployment. The working-age population, in line 4, is equally unimportant, for obvious reasons.

Line 5 describes the cyclical role of the labor-force participation rate. On this topic, Okun wrote “...the postwar record has convincingly delivered the verdict that a weak labor market depresses the size of the labor force. But the magnitude and time of the effect is not clear.” Many commentators hold the same view today. The results are clear that—in the useful framework Okun pioneered—participation is not at all cyclical. The coefficient is essentially zero with a small standard error. The confidence interval admits no meaningful cyclical behavior. Of course, a large unexpected reduction in participation occurred after the 2008 financial crisis. I take up the subject of how much of that experience can be related to the slack labor market later in this paper.

Line 6 shows the direct effect of unemployment operating through the log of the employment rate, which is one minus the unemployment rate, so the regression fits essentially perfectly and the standard error of the coefficient is effectively zero.

Line 7 shows that Okun was right that hours per worker was an important part of the cyclical relationship between labor-market tightness, measured by unemployment, and

output. A slacker labor market with higher unemployment results in a decline in hours per worker and thus a decline in output.

Line 8 deals with the bridge between the total economy, which contributes the data on the labor market, and the private economy, the subject of productivity measurement. Luckily the two economies move together over the cycle—the coefficient is essentially zero.

Finally, line 9 shows that labor quality has a small but statistically unambiguous role at cyclical frequencies. The positive coefficient supports the hypothesis that times of high unemployment are times of higher labor quality, because unemployment differentially strikes lower-skill workers.

3.2 Volatility of the cyclical and non-cyclical parts of the components

The third and fourth columns of Table 1 show the standard deviations of the implied cyclical and non-cyclical parts of each component, as percents at annual rates. The cyclical part is the coefficient in the first column multiplied by 4 times the standard deviation of the change in the unemployment rate in percents. The non-cyclical part is the total variable at annual rates less the cyclical part—the residual in the regression.

Line 1 shows that the changes in real GDP are evenly split in volatility between the cyclical part and the non-cyclical part, each with standard deviations of about 3.3 percent at annual rates. It is definitely not the case that real GDP evolves as a smooth growth path plus a cyclical part that tracks unemployment. As the rest of the table shows, there are quite a few sources of non-cyclical volatility of real GDP.

Figure 3 shows the implications of the decomposition of the movements of the level of the log real private GDP. The cyclical part tracks unemployment but is scaled by the estimated coefficient of -2.125 . A log-linear trend is removed, so the series starts and ends at the same value. The non-cyclical part has no trend, because it is the difference between the actual level of GDP and the fitted level.

In levels rather than first differences, the relative importance of the non-cyclical part of GDP is greater. The forces accounting for those movements tend not to revert to previous means, whereas the cyclical part is mean-reverting. The non-cyclical part grew rapidly between 1948 and 1978, then contracted a bit, grew more in the late 1990s and early 2000s, reaching a peak in 2002, then fell gradually until 2006, began to contract quickly, and even




Figure 3: Cyclical and Non-Cyclical Parts of Detrended Real Private GDP

more quickly after the crisis in 2008. The decline in the non-cyclical part continued at a rapid pace through the end of the sample in 2014.

Line 2 of Table 1 shows that the non-cyclical part of productivity growth is much larger than the cyclical part. Figure 4 confirms that the same is true in log-levels, using the same decomposition as for GDP. Non-cyclical movements of total factor productivity are a large part of the story of the movements of real GDP since 1948, both in terms of the non-cyclical part and the entirety of GDP.

Lines 3 and 4 of the table show that the slow-moving components, capital input and population, contribute moderate amount of non-cyclical movement to output, even though their cyclical contributions are essentially zero.

Line 5 shows the substantial contribution of non-cyclical movements of labor-force participation. Figure 5 reveals a bulge of rising participation starting in the mid-1960s and lasting until 2000, followed by a downward acceleration that worsened after the crisis but failed to reverse when the labor-market began to recover.

The tiny contribution of the employment rate to the non-cyclical part of output shown in line 6 arises entirely from the nonlinear relation between the log of one minus the unemployment rate and the unemployment rate itself. It arises because, following Okun, almost




Figure 4: Cyclical and Non-Cyclical Parts of Detrended Log Total Factor Productivity




Figure 5: Cyclical and Non-Cyclical Parts of Detrended Labor-Force participation




Figure 6: Cyclical and Non-Cyclical Parts of Detrended Hours per Worker

all research has treated the unemployment rate, rather than the log-employment rate, as the cyclical indicator.

Line 7 shows that non-cyclical movements of hours per worker are a substantial contributor to non-cyclical movements of output, with a standard deviation more than twice the cyclical contribution. Figure 6 shows that the low-frequency non-cyclical part of hours per work made important contributions to growth early in the period and depressed growth after 1970.

Line 8 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, line 9 of the table and Figure 7 show that labor quality was an important part of non-cyclical movements of through the late 1970s and a positive contribution until the mid-1990s.

3.3 Stability of the results

The relations between unemployment and the components estimated in the literature on Okun's Law are descriptive, not structural. There is no presumption of the stability over time expected of deep structural parameters. Nonetheless, as the top line of Table 2 shows, Okun's




Figure 7: Cyclical and Non-Cyclical Parts of Labor Quality

regression using data ending in 1960 gave very close to exactly the same basic coefficient as found here: Okun found that the expectation of real growth conditional on unemployment was -2.1 times the unemployment rate. Using the current data for his sample period, ending in 1960, the estimated coefficient is -2.175 , remarkably close to the estimate over the period ending in 2014, -2.125 . An empirical law is a stable relationship found repeatedly in the relevant data. By that standard, Okun definitely propounded a law.

The right-hand two columns of the table show estimates using data starting in 1984, the beginning of the era of stable monetary policy. The basic Okun coefficient is about -1.8 . The evidence that the decline is not just the result of sampling error is moderately strong, but the magnitude of the decline is not large.

The results for the components shown in lines 2 through 9 of the table show fairly small differences between the coefficients for Okun's sample period and the entire period. Some of the differences between the results for the sample beginning in 1984 and the results for the entire period are notable: the coefficient for productivity is only half as large, as is the coefficient for hours per worker.

<i>Line</i>	<i>Component</i>	<i>Full sample, 1948:2 to 2014:4</i>		<i>Okun's sample, 1948:2 to 1960:4</i>		<i>Recent sample, 1984:1 to 2014:4</i>	
		<i>Coeffi- cient</i>	<i>Standard error</i>	<i>Coeffi- cient</i>	<i>Standard error</i>	<i>Coeffi- cient</i>	<i>Standard error</i>
1	Private real GDP	-2.125	(0.128)	-2.175	(0.279)	-1.773	(0.191)
2	Total factor productivity	-0.911	(0.124)	-1.064	(0.252)	-0.474	(0.193)
3	Capital input	-0.032	(0.015)	-0.019	(0.017)	-0.105	(0.034)
4	Population 16 and over	0.018	(0.017)	0.025	(0.030)	-0.031	(0.031)
5	Labor-force participation rate	0.025	(0.033)	0.093	(0.078)	-0.056	(0.048)
6	Employment rate	-0.722	(0.001)	-0.718	(0.002)	-0.704	(0.002)
7	Hours per worker	-0.516	(0.074)	-0.741	(0.174)	-0.251	(0.103)
8	Ratio of private to total hours of work	-0.051	(0.056)	0.230	(0.147)	-0.343	(0.068)
9	Labor quality	0.063	(0.021)	0.019	(0.012)	0.192	(0.053)

Table 2: Results for Sub-Periods




Figure 8: Real Private GDP, 2000 to 2014

4 Events in the U.S. Economy between 2000 and 2014

The variables considered in this paper in connection with Okun's Law behaved unusually in recent years. The framework of Okun's Law provides a useful lens to study the anomalies. In addition to real GDP, five of the components had important movements of their non-cyclical parts between 2000 and 2014.

Figure 8 shows that the non-cyclical part of real GDP began to drop, relative to its normal trend, in 2006. Its decline was interrupted by a small increase in 2009 and 2010, when the cyclical part began to decline sharply. Then the non-cyclical part resumed declining. When the cyclical part began to rise, as unemployment subsided, continuing declines in the non-cyclical part fully offset the cyclical recovery. The total decline in the non-cyclical part from 2006 to 2014 was an astonishing 16 percent.

Answering the question, *Why Has the Unemployment Rate Fared Better than GDP Growth?*, is a matter of studying the sources of the striking decline in the non-cyclical part, which is simply actual log GDP less the part of log GDP that moves with unemployment. Five of the eight components moved non-trivially between 2000 and 2014.




Figure 9: Total Factor Productivity, 2000 to 2014

4.1 Productivity

Total factor productivity, adjusted for trend, fell about 8 percent from its peak in 2006 until 2014, half of the total decline over the same period in the non-cyclical part of real GDP. The cyclical part of productivity fell by five percent, reflecting the coefficient of 0.9 in the log-real-GDP relation to unemployment and a large increase in unemployment. But a good deal of that decline didn't happen in this contraction, so the non-cyclical part rose, because it is calculated as the difference between actual productivity and the estimated cyclical part. Once unemployment started falling, in 2011, the cyclical part was responsible for part of the continuing decline in the non-cyclical part.

TFP 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 10 shows a detrended 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 detrended in this paper.




Figure 10: Investment in Productivity Improvements

Figure 10 shows that IP investment grew faster than normal during the period of high TFP 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.

Recall the earlier discussion of how Fernald's productivity measure based on the Solow framework treats intellectual-property investment. The effect of the investment on productivity is the excess over the amount included in the capital stock. The collapse of IP investment was one of the sources of the noncyclical shortfall in the capital stock documented earlier and to be discussed further in the next subsection. The effect on total factor productivity is any extra effect not captured by the decline in the measured stock of intellectual property.

The relation between the Internet and productivity growth has received a good deal of attention. The Internet and related computer-based technical advances seemed to account for rapid productivity growth in the late 1990s and early 2000s. Advances such as social internet services and mobile internet apps continued in later years, but productivity growth diminished. Internet boosters have complained that the productivity accounting system fails

to include new Internet-based products and thus understates the growth of productivity. Syverson (2016) makes a wide-ranging investigation of the hypothesis that true TFP growth has exceeded the BLS's measure thanks to omission of improvements relating to the Internet and other sources in the past decade. His conclusion is uniformly negative. In particular, the scale of potentially omitted Internet-based products is much too small to account for the shortfall in TFP growth.

Fernald (2015) carries out a thorough investigation of the productivity slowdown. His basic conclusion is that 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. He writes, "By the middle of the first decade of the twenty-first century, the low-hanging fruit of IT had been plucked." He investigates the potential role of other influences, including the housing boom, the labor-market slack that followed the financial crisis, and declining rates of new-business creation, but concludes against assigning them much of a role. He observes that the productivity slowdown began several years before the crisis. His general conclusion is that higher productivity growth around 2000 was special and that lower productivity growth prior to the mid-1990s was normal and the best that can be expected now.

Gordon (2014) enumerates technologies that have been contributing to productivity growth in the past decade and promise further contributions in coming decades. He finds that none compares to the contributions in similar areas during the most recent sustained period of high productivity growth, which ended in 1972. He finds that "Pharmaceutical research has reached a brick wall of rapidly increasing costs and declining benefits..." and that "[Robots] can think but can't walk, or they can walk but can't think." He makes the good point that, as the price of computing power approaches zero, its marginal benefit also approaches zero; computing saturation occurs and the contribution of computing to productivity growth must end. Driverless cars offer hardly any advantage over existing cars, as people might as well drive once they occupy a car.¹

¹Reading this section of the paper reminded me of a speech I attended in 2005 by the CEO of Qualcomm, the company that makes the chips that power cellphones. His theme was that there was a coming crisis in his business because cellphone makers couldn't think of enough different functions to put on a phone to keep the most recent chips fully occupied. In 2007, Apple introduced the iPhone and chips have been fully occupied ever since.




Figure 11: Capital Input, 2000 to 2014

4.2 Capital input

Figure 11 shows the decline in capital input that began shortly after the 2008 crisis. It is responsible for about four percentage points of the shortfall of non-cyclical real GDP. Fernald's measure of capital input includes the services of plant, equipment, inventories, and R&D and other intangibles.

Real business investment in equipment is a major fraction of capital formation. It embodies many of the new technologies that account for productivity growth. Figure 12 shows its values detrended as for other indexes. 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 in half in 2009, a much larger percentage drop than in any previous recession in the years since 1948. In the recovery, it has returned to trend, well below its level in the previous two decades (detrended).

An important determinant of business investment is the payoff to owners of capital. A potential explanation for the extraordinary weakness of investment following the financial crisis would be a finding that capital was not earning as much as in normal times. But, as Figure 13 shows, the earnings of capital, measured as the sum of business profits, interest




Figure 12: Equipment Investment

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.

Investment in plant, equipment, and IP has remained weak at the same time that investment in job creation has returned to normal. The puzzle of low investment has yet to be solved.

4.3 Working-age population

As Figure 14 shows, the working-age population declined sufficiently to cut real GDP by 2.5 percent relative to its trend, between 2008 and 2014. The decline was the result of (1) smaller cohorts of young people passing their 16th birthdays, (2) lower net immigration, and (3) higher mortality rates.

4.4 Labor-force participation rate

Figure 15 shows that a decline of participation in the labor force accounted for a decline of about three percent in real GDP after 2008, relative to trend. This decline is non-cyclical, as the earlier regression found essentially no systematic relation between unemployment and participation (contrary to the beliefs of Okun and many macroeconomists today). Prior to




Figure 13: Business Earnings as a Ratio to the Value of Capital




Figure 14: Working-Age Population, 2000 to 2014




Figure 15: Labor-Force Participation Rate, 2000 to 2014

the financial crisis in 2008, 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 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 16 shows the percentages of the populations who are in the labor force. There is no detrending in the figure.

Figure 16 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 force specific to the post-crisis years depressed participation.

Table 3 provides some 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 2007, years when the labor market was unaffected by the crisis, small declines in participation, averaging 0.8 percentage points, occurred in all four cate-




Figure 16: Labor-Force Participation Rates

gories of family income. Between 2007 and 2013, participation *rose* among members of the poorest quarter of families, fell just a bit in families in the second quartile, and fell by 2.5 percentage points in the upper half, the third and fourth quartiles. Essentially all of 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.

Labor force participation among prime-age workers across household income distributions

	2004	2007	2013
Total	83.8%	83.0%	81.2%
1 st quartile (lowest income)	62.3%	61.2%	61.5%
2 nd quartile	80.0%	78.0%	77.6%
3 rd quartile	88.0%	87.3%	84.8%
4 th quartile (highest income)	91.9%	91.4%	89.9%

Table 3: Role of Family Income

	<i>Personal care, including sleep</i>	<i>Market work</i>	<i>Education</i>	<i>Leisure</i>	<i>Other</i>
Men	1.3	-1.6	-0.1	1.6	-1.2
Women	2.2	-1.4	0.0	1.2	-2.0

Table 4: Changes in Weekly Hours of Time Use, 2007 to 2014, People 15 and Older

Table 4 investigates how people spent the time freed up by reduced work and job search. It compares time allocations in 2014 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 post-crisis decline in participation 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 1, but for the period from 2008, first quarter, to 2016, first quarter, 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 LFPR” (p. 11). I 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 LAUS data for the post-crisis period is not typical of the data for earlier periods. Over the entire duration of the LAUS, 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.

Elsby, Hobijn and Şahin (2015) are skeptical of the value of studying the joint behavior of the participation rate and the unemployment rate. They agree that there has been little cyclical movement in participation historically, and they do not focus on the unusual post-crisis decline in participation. Rather, they believe that there are aspects of participation that are not captured by aggregate participation but are measured as the gross flows among participation, unemployment, and employment. I find their case against studying the level of participation in the population unconvincing, though I agree that there is much to learn from the flows. Integrating flows into models of individual labor-market behavior is currently an active area of research.

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.

4.5 Hours per worker

Although hours per worker is a cyclical variable, it has a noticeable non-cyclical part, as shown Figure 17. And that part was the only favorable contribution found in this paper to the non-cyclical movement of real GDP during the post-crisis period.




Figure 17: Hours per Worker, 2000 to 2014

5 Concluding Remarks

Table 5 gives this paper’s basic answers to the question posed in its title: Why did unemployment fare better than GDP? It shows the mostly bad news from sources of changes in real GDP over the period from the fourth quarter of 2007 to the fourth quarter of 2014 that are not the normal results of the movements in cyclical resource utilization as measured by unemployment. The contributions are stated as changes in the log of real GDP interpreted as percent changes. They measure the difference difference between what would have happened if real GDP had tracked its normal relationship to GDP, as expressed in Okun’s Law, and what actually happened to real GDP.

The three big negative components of the breakdown are productivity, capital input, and labor-force participation, each contributing about four percentage points to the shortfall of real GDP. The decline in the working-age population contributed another two percent. The favorable components, coming in here with negative signs, are hours per worker, the ratio of private to total hours of work, and labor quality. These saved the big decline in real GDP relative to trend from being 2.4 percentage points worse.

These results illustrate the great importance for macroeconomics of forces affecting output and real income apart from the traditional focus on demand-driven fluctuations in resource

<i>Component</i>	<i>Shortfall, percent</i>
Total factor productivity	4.4
Capital input	4.0
Population 16 and over	2.2
Labor-force participation rate	3.6
Hours per worker	-1.7
Ratio of private to total hours of work	-0.1
Labor quality	-0.6
Private real GDP	11.7

Table 5: Why Unemployment Fared Better than GDP, 2007 to 2014

utilization. The idea that the evolution of real GDP follows short-run fluctuations in demand and a smooth underlying trend was shown to be completely wrong in the years following the crisis, and not for the first time. Large fluctuations in the non-cyclical components of GDP occurred at many other times in the years since 1948 studied in this paper.

6 Related Literature and Sources

I will not attempt to review systematically the large and growing literature on the macroeconomics of the financial crisis and ensuing slump and stagnation. Many references appear in my chapter, Hall (2016), for the forthcoming Volume 2 of the *Handbook of Macroeconomics*, and many of the other chapters in the volume treat the subject.

Table 3 is taken from Hall and Petrosky-Nadeau (2016).

See the spreadsheet available at Stanford.edu/~rehall for complete sources and calculations for the tables and figures in this paper.

References

- Elsby, Michael W.L., Bart Hobijn, and Ayşegül Şahin, “On the Importance of the Participation 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.
- Fernald, John G., “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” September 2012. 2012-19, Federal Reserve Bank of San Francisco. Updated regularly on Fernald’s FRBSF website.
- , “Productivity and Potential Output before, during, and after the Great Recession,” *NBER Macroeconomics Annual*, 2015, 29 (1), 1–51.
- Gordon, Robert J., “The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections,” Working Paper 19895, National Bureau of Economic Research, February 2014.
- 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.
- Okun, Arthur M., “Potential GNP: Its Measurement and Significance,” *Proceedings of the Business and Economics Statistics Section, American Statistical Association*, 1962, pp. 98–104.
- Polivka, Anne E. and Stephen M. Miller, “The CPS After the Redesign: Refocusing the Economic Lens,” March 1995. Bureau of Labor Statistics.
- Solow, Robert M., “Technical Change and the Aggregate Production Function,” *The Review of Economics and Statistics*, 1957, 39 (3), 312–320.

Syverson, Chad, “Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown,” Working Paper 21974, National Bureau of Economic Research, February 2016.