

WHAT RECOVERY?

The Case for Continued Expansionary Policy at the Fed

REPORT BY **J.W. MASON**
JULY 25, 2017



ABOUT THE AUTHOR

J.W. Mason is a Fellow at the Roosevelt Institute, where he works on the Financialization Project, and an assistant professor of economics at John Jay College, CUNY. His current research focuses on the history and political economy of credit, including the evolution of household debt and changing role of financial markets in business investment. He also works on history of economic thought, particularly the development of macroeconomics over the twentieth century.

ACKNOWLEDGMENTS

We thank Laurence Ball and Jason Furman for their comments and insight. Special thanks for the research assistance from Amanda Page-Hoongrajok and John Sturm. Roosevelt staff Nellie Abernathy, Amy Chen, Mike Konczal, Katy Milani, Marshall Steinbaum and Alex Tucciarone all contributed to the project.

This report was made possible with generous support from the Open Philanthropy Project.

Executive Summary

This paper critically examines the widely held view that the U.S. economy is today operating at close to potential. The paper makes five core arguments. First, GDP remains well below both the long-run trend and the level predicted by forecasters a decade ago. In 2016, real per capita GDP was 10 percent below the Congressional Budget Office's (CBO) 2006 forecast, and shows no signs of returning to the predicted level. Second, the fall in the employment-population ratio over the past decade is not due to demographics. At most, 40 percent can be attributed to the aging population, while the remainder represents declining employment rates within age groups, particularly younger ones. Third, the exceptionally slow productivity growth of the past decade can be understood as the result of weak demand, including the equally exceptionally weak growth of investment spending. Fourth, while individual pieces of the macroeconomic picture might seem to have supply-side explanations, the overall picture is difficult to square with a story of exogenous fall in the economy's productive potential. This combination of weak output growth, low inflation, slow wage growth and low interest rates is a textbook result of a negative demand shock, but is much harder to square with a fall in supply. And fifth, under current conditions, the costs of underestimating potential output are much greater than the costs of overestimating it. Under current conditions, higher inflation and a rising wage share are arguably desirable; or at least should be considered much less costly than the stagnant incomes, underemployment and waste of productive resources that will result from underestimating potential output.

Table of Contents

Introduction and overview	6
Potential output may have different meanings when we consider shorter or longer periods.	7
Changes in short run potential output as a result of demand are increasingly recognized in the form of “hysteresis.”	9
Operationalizing potential: Full employment and price stability	11
Outline of the paper	13
Section One: What recovery?	18
1.1 U.S. GDP remains at about 15 percent below the pre-recession trend.	22
1.2 Most of the slowdown was not anticipated by forecasters. The declining output gap reflects estimates of potential GDP converging to actual GDP, rather than actual GDP returning to trend.	24
1.3 The decline in growth reflects a mix of lower labor force participation and slower productivity growth.	27
Section Two: Labor force participation	35
Section Three: Productivity and demand	45
3.1 There are many reasons to think that demand could affect productivity.	45
3.2 The slowdown productivity growth has coincided with exceptionally weak investment spending.	50
3.3 The productivity growth slowdown is widely shared across sectors, even less technology-intensive ones.	53
3.4 Productivity and the business cycle	56
Section Four: The big picture: Overheating?	60
4.1 The overall state of the economy is hard to reconcile with a story in which we are facing supply constraints.	60
4.2 Potential is not a hard line; policy needs to take into account not only where it is, but the consequences of guessing wrong.	66
Conclusions	74

Introduction and overview

A central question guiding macroeconomic policy is how close the economy is to potential, or the maximum amount that can be produced given the existing labor force, real resources and productive capabilities.

When output is below potential, the economy is demand-constrained: spending more money—whether by households, businesses or government—will result in more people being employed, and more goods and services being produced. But once potential is surpassed, supply constraints become more important—output cannot be raised further unless the real productive capacity of the economy is increased. As a result, attempts to further raise spending will be increasingly stymied by bottlenecks and shortages, and will result only in “overheating” and inflation, along with a redistribution of income toward the scarce factors of production. More output and employment are almost always considered good things by policymakers, while inflation is invariably considered a bad thing.¹ The task of macroeconomic policy is to keep the level of aggregate spending, or demand, just in line with the economy’s potential output. Concretely: When the Federal Reserve (the Fed) thinks that gross domestic product (GDP) is well below potential, it lowers interest rates or takes other, unconventional actions to increase the supply of credit. When it thinks the economy is above potential, it raises rates or otherwise seeks to rein credit in.

If potential output could be directly measured, macroeconomic policy would still involve challenges, but at least it would not be hard to agree on the general direction of policy. Unfortunately, there is no way to measure potential, or even a clear standard for how it should be measured in principle. In practice, it is defined by a number of different measures and concepts. While these measures are sometimes in broad agreement, they sometimes conflict, and therefore can be defined and interpreted in different ways. The result is that at times, like today, the state of the economy may look very different to different observers. Where some see an economy in serious danger of overheating and in need of a firm application of the brakes, others see an economy that has barely recovered from the Great Recession and remains far from its full potential. Right now this debate is most pressing at the Fed, whose leadership must decide whether to raise interest rates and unwind the

Demand: Demand is defined as the aggregate spending in an economy, meaning all spending by households, businesses and government. An increase in demand means higher desired spending at a given level of income.

¹ We will briefly return in the final section to the question of whether this fear of inflation is justified, at least in present conditions.

measures it took to stabilize the financial system in 2008-2009. But it should be a concern for anyone involved in macroeconomic policy debates. A whole range of policies—from tax cuts to Medicaid expansion—look quite different in an economy that needs more demand, than in one that is already running at full capacity.

Unfortunately, there is no way to measure potential, or even a clear standard for how it should be measured in principle ... Where some see an economy in serious danger of overheating and in need of a firm application of the brakes, others see an economy that has barely recovered from the Great Recession and remains far from its full potential.

POTENTIAL OUTPUT MAY HAVE DIFFERENT MEANINGS WHEN WE CONSIDER SHORTER OR LONGER PERIODS.

For our purposes, it is important to distinguish between three concepts of potential, which we will call short run, medium run and long run. Short run potential is the level of output consistent with low, stable inflation. Once spending rises above this level, we expect to see inflation that is rising and/or above the level desired by policymakers. Conceptually, this is the level at which attempts by households, businesses and governments to buy more begin to encounter significant supply constraints. Short run potential corresponds to a normal level of measured unemployment, though not to “full employment” in the sense that everyone who is willing and able to work has a job. Medium term potential is the maximum level of output that will eventually be reached simply by increased demand, or spending. The distinction matters if we believe that the economy is characterized by hysteresis—that is, by adjustment of productive capacity to demand conditions. In a world without hysteresis, short run and medium run inflation would be the same: once supply constraints are reached, additional spending can only drive up prices, not output. But in a world with hysteresis, higher demand may eventually translate into greater productive capacity. A strong labor market with rising wages will draw more people into the workforce, and tight labor and product markets will encourage businesses to shift toward more productive, capital-intensive products and techniques. In this case, higher demand may initially lead to rising inflation, but the higher inflation will be temporary as increased demand eventually

leads to increased supply. Finally, long run potential can be thought of as the maximum level of output consistent with the underlying technical facts of population, technology, natural resources and so on. Reaching this higher level of output may require additional policy interventions beyond simply increased expenditures. These three definitions are often not distinguished in practice. In particular, most of the widely used macroeconomic models assume that the economy's productive capacity develops independently of the level of spending. In that case, there would be no difference between what we are calling short run and medium run potential. But in the current setting, where hysteresis is important, it is crucial to distinguish them. If increased demand does produce higher inflation, that does not mean that it can no longer raise output or that continued high demand would lead to higher inflation indefinitely.

The goal of macroeconomic policy is not to avoid inflation at any cost, but rather to strike the best possible balance between rising living standards, full employment, price stability and other objectives.

In this paper, we are concerned with medium term potential. We aim to answer the question of whether there is space for increased demand to raise output. Throughout the paper, the term “potential” is, unless otherwise noted, used in this sense. It is possible, though not certain, that continued strong demand will soon lead inflation to rise over 2 percent—that is, the economy is indeed close to reaching its short run potential. As we argue in Section 4.2, higher inflation is not necessarily a bad thing for the U.S. at this moment. But in any case, this should not be the focus of debate. The goal of macroeconomic policy is not to avoid inflation at any cost, but rather to strike the best possible balance between rising living standards, full employment, price stability and other objectives. The critical question for policymakers today—at the Fed in particular—is whether strong demand could substantially raise output and employment. If the answer is affirmative, there is a strong case for sustaining expansionary policy, whether or not there is some higher inflation along the way.

CHANGES IN SHORT RUN POTENTIAL OUTPUT AS A RESULT OF DEMAND ARE INCREASINGLY RECOGNIZED IN THE FORM OF “HYSTERESIS.”

The idea that recessions, or weak demand in general, can have lasting effects on the economy’s productive potential is often referred to as “hysteresis.” The textbook view sees demand and supply evolving independently—demand depends on households and businesses’ decisions about how much to spend out of their income, a consideration influenced by factors like interest rates; while supply depends on factors like labor force growth and technological change, which are not affected by current spending decisions. Advocates of hysteresis argue that the two can’t be so neatly separated—changes in spending can have lasting effects on the economy’s productive capacities. Hysteresis may explain the failure of output to return to its earlier trend. As several economists at the New York Fed recently put it, “One interpretation of the listless recovery is that recessions inflict permanent damage on an economy’s productive capacity. For example, extended periods of high unemployment can lead to skill losses among workers, reducing human capital and lowering future output” (Acharya et al. 2016).

Hysteresis may explain the failure of output to return to its earlier trend. As several economists at the New York Fed recently put it, “One interpretation of the listless recovery is that recessions inflict permanent damage on an economy’s productive capacity. For example, extended periods of high unemployment can lead to skill losses among workers, reducing human capital and lowering future output.”

Hysteresis has clear implications for policy: It makes any shortfall of output from potential more costly than the reduced production in just that year:

“To the extent that hysteresis is present, it implies that deviations in output from its optimal level are much longer lasting and thus more costly than usually assumed. The implication is straightforward, namely that monetary policy should react more strongly to output movements, relative to inflation” (Blanchard, Cerutti and Summers 2015).

Or as the New York Fed economists put it, “Monetary policy rules that fare well in normal times can lead to pathological outcomes when faced with large shocks.” This is because if the authorities treat the damage done by the recession as a sign that the economy’s capacity is simply lower than believed, they will refuse to allow the strong expansion that could reverse the damage. In effect, a belief that hysteresis just reflects the “new normal” can be self-confirming. Instead, faced with hysteresis, the central bank should aim for a period of exceptionally strong demand, with levels of unemployment and wage growth that would normally be considered inflationary: “Just as recessions damage potential output, booms can repair it.”

Faced with hysteresis, the central bank should aim for a period of exceptionally strong demand, with levels of unemployment and wage growth that would normally be considered inflationary: “Just as recessions damage potential output, booms can repair it.”

Similar conclusions have been reached by a number of other economists who have studied the long-term effects of recessions, especially in Europe, where long-term growth rates have not been as stable as in the U.S. Laurence Ball, for instance, finds that a temporary period of elevated inflation can permanently raise GDP—exactly the opposite of the conventional view (Ball 2009). Along the same lines, DeLong and Summers (2012) show that if demand shortfalls have even modest persistent effects, expansionary fiscal policy can pay for itself via the higher growth it generates.

The most common version of the hysteresis story focuses on the labor market. In these cases, prolonged unemployment reduces workers’ ability and/or motivation to find a job. In Ball’s version, for instance, the main effect of hysteresis is to raise the “natural” or “non-accelerating rate of inflation” rate of unemployment—the minimum unemployment rate at which inflation is stable (Ball 2014). The idea is that after a deep recession, we may see rapid wage growth even while unemployment is still high, because the long-term unemployed are not easily employable, and so businesses must compete for the small pool of recently unemployed. If central banks see this as a reason to raise rates, they will in effect lock in the new, higher level of unemployment. But if demand remains strong, eventually the long-term unemployed will reconnect with the labor market and inflation will stabilize at a lower level of joblessness.

This version of the argument is very relevant for Europe, where the unemployment rates regarded as “natural” by the authorities are now much higher than before the recession: 10 percent, up from 7 percent in Italy; 17 percent, up from 10 percent in Greece; and 21 percent, up from 12 percent in Spain. It is hard to believe that a fifth of the Spanish workforce now

needs to be unemployed for prices to remain stable; it is much more natural to suppose that this apparent upward shift reflects disruptions to the labor market as a result of the country's financial crisis and deep recession, which a period of strong growth could reverse.

In the U.S., however, measured unemployment is now close to its lowest point in the 2000s, making this version of the hysteresis argument inapplicable. It is true that labor force participation has declined since the recession, which could reflect similar factors. But even this explains only half of the deviation of GDP from its previous trend. The majority of the shortfall reflects a decline in output per worker, which doesn't fit the standard hysteresis story. So when thinking about hysteresis in the current U.S. context, we need to consider a broader set of factors, as suggested by DeLong and Summers:

“A depressed economy is one in which many workers are without employment for an extended period. As a consequence, many see their skills, the networks they use to match themselves with vacancies in the labor market and their morale all decay. A depressed economy is also one in which investment is low, the capital stock is growing slowly if at all, and entrepreneurial exploration is low, and it is certainly possible that this deficit is not made up quickly” (DeLong and Summers 2012).

If a period of weak demand erodes skills and discourages investment and innovation, presumably a period of strong demand would do the opposite. Insofar as slow productivity growth reflects, at least in part, the continuing effect of the recession, as suggested by DeLong and Summers, there is more space for expansionary policy than conventional measures of potential would imply.

If a period of weak demand erodes skills and discourages investment and innovation, presumably a period of strong demand would do the opposite.

OPERATIONALIZING POTENTIAL: FULL EMPLOYMENT AND PRICE STABILITY

There has been a great deal of debate in recent years about which policy instruments should be used to close an output gap. This is a significant step forward from a decade ago, when the overwhelming consensus was that adjustments to a policy rate by a central bank could and should be the sole tool used to correct macroeconomic imbalances. But while there has been a new interest in unconventional monetary policy, and renewed interest in the use of fiscal policy to close output gaps, relatively less attention has been paid to the measurement and meaning of output gaps themselves.

While there has been a new interest in unconventional monetary policy, and renewed interest in the use of fiscal policy to close output gaps, relatively less attention has been paid to the measurement and meaning of output gaps themselves.

Potential output may be defined as the maximum output achievable given the real resources (labor, technology, natural resources) available to the economy, or as the level of output consistent with stable inflation, the level of output corresponding to full employment, or the level of output to which the economy tends to converge over the long term as business cycle fluctuations average out.² These different definitions may be equivalent in some economic models, and they may sometimes lead to similar practical assessments. But conceptually, they are distinct, and imply different approaches to measuring the gap between current and potential output.

For policy purposes, the two most important measures of whether output is at potential are full employment and price stability.

For policy purposes, the two most important measures of whether output is at potential are full employment and price stability. In theory, these measures are supposed to be equivalent—full employment is the minimum level of unemployment achievable without inflation accelerating past the central bank’s speed limit. In practice, inflation does often rise at the same time that unemployment falls, but not always, and not in strict proportion. The central bank must therefore consider both measures independently—a “dual mandate” that, in the U.S., is written into law.

The Fed defines price stability as 2 percent inflation. (Its preferred inflation measure is the Personal Consumption Expenditure deflator, or PCE, and not the more familiar Consumer Price Index, or CPI.) This 2 percent target, while widely shared by central banks around the world, is a relatively recent development. Until the 1990s, most central banks, including the Fed, pursued low inflation but did not set an explicit numerical target (Krugman

² Potential output is often estimated as a smoothed version of actual GDP, using a Kalman, Hedrick-Prescott or other statistical filter. While this approach is conceptually straightforward, the actual results can be quite sensitive to the filter used and the choice of parameters (Akin, Kose and Ohnsorge 2014).

2014). Full employment is a fuzzier target; there is no explicit numerical definition, and estimates of the “natural” or “non-accelerating inflation” unemployment rate have varied widely over the years. But while there is no explicit target, members of the Federal Open Markets Committee (FOMC), the Fed’s governing body, do release estimates of the long-run unemployment rate toward which they expect the economy to converge. As of the most recent FOMC meeting, this ranges from 4.5 percent to 5 percent. This is a reasonable guide to what the current Fed leadership sees as full employment.

By these measures, as of the first quarter of 2017, the U.S. economy appears close to potential, if perhaps not quite there yet. As of January, annual PCE inflation stood at 1.9 percent, up from an average of 1.4 percent in 2016. This is the closest that inflation has gotten to the Fed’s target in almost five years. This was in part due to rising food and energy prices. With these volatile components stripped out, “core” inflation was only 1.7 percent.

As for unemployment, the headline unemployment rate had fallen to 4.7 percent, near the middle of the range considered by FOMC members as the normal, full-employment level. This is a fall of more than half from the 10 percent reached in October 2009, at the bottom of the Great Recession. Since then it has gradually fallen, reaching 5 percent in September 2015, and now 4.7 percent. This is still a bit above the 4.6 percent average unemployment for the year before the Great Recession, and a bit above the 4 percent that unemployment averaged in the year before the 2001 Recession. But by the headline measure, unemployment is clearly close to, if not quite at, the level considered by the Fed to be full employment.

While most economists continue to regard the pre-recession targets of 2 percent inflation and around 5 percent unemployment as appropriate goals for policy, a few have argued that under today’s conditions, significantly more expansionary policy is needed. For example, Laurence Ball argues, “The Fed should seek to push the unemployment rate well below 5 percent, at least temporarily. A likely side effect would be a temporary rise in inflation above the Fed’s target, but that outcome is acceptable” (Ball 2015). We will return to this question in Section 3.

OUTLINE OF THE PAPER

In this paper, we argue that the balance of evidence points to substantial additional space for expansionary policy. There is good reason to think that the economy does not yet face hard supply constraints—that increased spending could sustainably raise employment and output. Fully utilizing productive potential, we argue, requires a high-pressure economy—one in which aggregate demand is allowed to continued rising past what is currently seen as the full employment level.

The paper is organized as follows:

In the **first section**, we show that the recovery since 2009 is, in a sense, a statistical illusion. Output remains about 15 percent below the pre-recession trend, a larger gap than at the bottom of the recession. The apparent closing of the output gap is entirely due to downward revisions to estimated potential. As far as GDP is concerned, the ball has not moved forward, the goal has moved backward. This 15-point gap is about equally due to the declining share of the population employed, and to slower growth in output per worker.

The recovery since 2009 is, in a sense, a statistical illusion. Output remains about 15 percent below the pre-recession trend, a larger gap than at the bottom of the recession. The apparent closing of the output gap is entirely due to downward revisions to estimated potential. As far as GDP is concerned, the ball has not moved forward, the goal has moved backward.

In the **second section**, we challenge the widely held view that the decline in labor force participation is simply the result of an aging population. The decline in labor force participation cannot be directly explained by demographics.

In the **third section**, we suggest that slower productivity growth is most naturally explained by weak demand. While productivity is often treated as fully determined by an economy's "endowments" of real resources and technology, there are many good reasons to expect strong demand to cause increased labor productivity. A number of lines of evidence support this interpretation of the productivity slowdown of the past decade, including the fact that the vast majority of industries have experienced slower productivity growth.

Finally, in the **fourth section**, we ask whether the economy as a whole looks like a case of overheating. On the contrary, we argue, while a myopic focus on one or another data series might support a story of binding supply constraints, the behavior of the economy as a whole is much more consistent with a situation of depressed demand—an extended recession. The overall picture also makes it unclear what actual danger is posed by overheating in the conventional sense. Most of the obvious costs of overheating—higher inflation; higher interest rates; a rising wage share—would be desirable under current circumstances.

The behavior of the economy as a whole is much more consistent with a situation of depressed demand—an extended recession. The overall picture also makes it unclear what actual danger is posed by overheating in the conventional sense. Most of the obvious costs of overheating—higher inflation; higher interest rates; a rising wage share—would be desirable under current circumstances.

Accepting that the economy has more room for expansionary policy does not, of course, imply endorsement for any particular route to increase spending. Higher public spending on foreign wars, immigration enforcement, environmentally destructive infrastructure projects, higher private spending by the wealthiest as a result of regressive tax cuts, or higher net exports as a result of an aggressive mercantilist trade policy would all boost aggregate demand. All, however, would also have grave social costs. But insofar as there is socially desirable spending that the government could undertake or encourage, we should not be deterred by the idea that the economy is already operating at maximum potential. Going forward, the problem of how to boost aggregate spending while meeting human needs should remain at the center of the policy agenda.

Insofar as there is socially desirable spending that the government could undertake or encourage, we should not be deterred by the idea that the economy is already operating at maximum potential. Going forward, the problem of how to boost aggregate spending while meeting human needs should remain at the center of the policy agenda.

IS HIGHER GDP NECESSARILY A GOOD THING?

While most economists and policymakers take it for granted that faster output growth is desirable, not everyone agrees. And indeed, there are valid reasons to be skeptical about the place growth occupies as the unquestioned goal of economic policy.

First, environmental resource constraints are real. To be sure, GDP measures the value placed on market purchases, and not any physical quantity. There is no necessary connection between the money payment associated with a given human activity and the demands it places on natural systems' capacity to reproduce themselves. In principle, a given throughput of physical inputs can be associated with an arbitrarily high GDP. But in our economies as currently organized, there is a close connection between rising GDP and increased use of fossil fuels. This raises obvious concerns about the consequences of faster growth for climate change. It also means that a spike in oil prices may cut off faster growth.

Second, the gains of higher growth are not shared equally across the income distribution. The U.S. labor market is highly segmented, with certain groups—based on occupation, credentials, race, gender, geography or other factors—benefitting more from strong demand than others.

It is possible that faster growth could encounter significant supply constraints (for example, rapid wage increases for

workers with scarce skills) well before the gains reach less favored segments. Profit income is also more cyclical than wage income, so, in the short term, an acceleration of growth is likely to be associated with a shift in income toward capital owners. If we take these concerns seriously, they suggest that the focus should be on direct job creation programs for people left out of the private market, rather than policies to raise output across the board.

A third reason to be skeptical about growth as a political project is that higher GDP does not necessarily improve people's wellbeing. There is no question that an increase of per-capita GDP from, say, \$300 to \$3,000 is almost invariably associated with a large improvement in welfare. This view is supported by a great deal of data on life expectancy, infant mortality, literacy and other direct measures of wellbeing. But the case is less clear-cut for a country that is already as rich as the United States. Both measures of human development and direct, survey-based measures of happiness suggest that by the time a country reaches the levels of per-capita income of the contemporary United States, the direct effects of further growth on wellbeing are weak and unreliable (Stiglitz et al. 2010).³

Wilkinson and Pickett

evidence, see Alexander (2012).

Not everyone will accept these three arguments against the primacy of growth, of course. But we believe that even if one does not accept them, there are other, equally strong considerations that point toward a need for expansionary policy. We believe there is a progressive case for growth even if we accept a strong version of the claim that higher GDP does not directly improve welfare.

First, there is good reason to believe that, while the gains of growth are not spread evenly, a tight labor market has a strong tendency to raise wages, flatten the income distribution and increase the social power of working people more broadly. Over the short run, it is true, profits respond more strongly than wages to growth. But sustained high growth and low unemployment almost always pulls up wages disproportionately, especially at the bottom of the distribution. This was demonstrated in the U.S. most recently in the 1990s. This period saw both sustained above-trend GDP growth and below-trend unemployment, and the only sustained period of a flattening income distribution and rising wage share in the past 40 years. On a larger scale, the same principle was illustrated in the 1960s and 1970s when sustained low unemployment in the U.S., Western Europe and Japan was associated with a large shift of income and, just as importantly, political power and status, toward wage earners. There is a strong argument that the mass social movements of this period, even such apparently non-economic ones as feminism and gay

rights, also got an important impetus from sustained labor market tightness.⁴

Second, the idea of a hard ceiling to potential output is an important part of the logic of scarcity that constrains the political imagination in much of the world today. If spending on public investment and income redistribution are good things on merit, the question becomes whether we can “afford” them. One good reason to think we can is that if the economy has unused productive potential, this applies in particular to the question of environmental constraints. While reducing carbon emissions may, in the long run, require curtailing certain forms of consumption (air travel, for instance), in the short run it will require a major increase in public and private investment. If the economy is still far from potential, that strengthens the case for increased spending on green energy, building retrofits, smart grids, new energy-storage technology, less fuel-intensive land use patterns and so on.

Finally, low-productivity jobs and underemployment are generally less intrinsically rewarding, more tedious and often more physically taxing, less secure and enjoy less respect and status than productive jobs. So an argument that the economy is still far from potential is in part also an argument for better jobs for ordinary Americans.

The classic discussion of the links between full employment and the new social movements of the 1960s is Armstrong, Glyn and Harris (1991). For the US, see also Jefferson Cowie (2010).

What recovery?

In this section we argue that, by historical standards, the economy has not recovered from the 2008 financial crisis in terms of GDP. It is true that measured unemployment has returned to the level considered equivalent to full employment, and price inflation is close to—if still short of—the Fed’s 2 percent target. But output remains a full 15 percent below the pre-2007 trend line, a gap that is getting wider, not narrower, over time. This is shown in Figure 1, which compares the actual path of real per-capita GDP to a linear trend from 1947 to 2007.

PER CAPITA GDP AND LONG-RUN TREND




FIGURE 1 Real per-capita GDP and 1947-2007 trend. The figure shows that between 1947 and 2007, real per-capita GDP showed a stable long-run growth rate, but since 2007 it has fallen below this trend. **Note:** The scale is logarithmic, so a constant rate of increase appears as a straight line. *Source: BEA, author’s analysis.*

As Figure 1 shows, there is no precedent for GDP’s current divergence from trend in postwar U.S. history. We will make three points about this output shortfall: First, the apparent closing of the output gap is due to downward revisions of estimated potential output. Second, most of this slowdown in potential output growth was not predicted prior to the Great Recession. And third, it is composed about equally of a decline in employment relative to population, and slower productivity growth. We consider each of these facts in more detail in the following sections, before returning to the economy as a whole in the final section.

First, a technical note: Some economists object to comparisons like that of Figure 1 on the grounds that we should not expect real per-capita GDP to follow a linear trend. That is, rather than expecting a steady long-term growth rate, we should expect the effects of booms and slumps to last indefinitely. While this is an important debate, we think the long-run trend is a meaningful point of comparison, for three reasons. First, it is the case that real per-capita GDP in the U.S. closely followed a linear trend in the 60 years after World War II, with periods of above-average growth followed by periods of below-average growth, and vice versa. (This is not necessarily true for other countries or periods.) Second, a strong tendency for GDP to return to trend is the default assumption in most practical policy debates, even if it is more controversial among academic economists. The assumption that most departures from the long-run trend are temporary is a feature of most forecasting models used by government and business. Third, and most fundamentally, it is hard to tell a story in which aggregate demand and supply matter if deviations from trend are persistent. The most prominent deviations from trend are recessions and recoveries, which nearly everyone—both economists and people in the policy and business worlds—agree are driven by shifts in aggregate demand—that is, by the flow of money through the economy.

To argue that these deviations are normally permanent, you have to take one of two radical positions. In one, you believe that the economy is always operating at potential—that recessions are due to technological regress or other declines in the productive capability of the economy, throwing out almost everything we know about real-world business cycles. “Real business cycle” economists, who deny that aggregate demand matters even in the short run, take this approach. But they are a distinct minority. With the other position, you may believe that changes in desired spending, or in the flow of money and credit, have permanent effects on output. This is a view that even fewer economists would accept. But without taking one of these two positions, it is hard to argue that economic growth doesn’t follow a more or less linear trend, even in cases where the statistical evidence is more ambiguous than for the postwar U.S.⁵ Of course, to the extent that demand-induced shifts in output persist in the long run, the case for expansionary policy is that much stronger. If even a small part of the effect of fiscal stimulus is permanent, such stimulus may pay for itself (DeLong and Summers 2012).

In 2006, the CBO was forecasting cumulative 16 percent growth in real, per-capita GDP over the next decade. Instead, growth has been just 6 percent. If GDP had grown as predicted, an additional \$1.7 trillion would have been produced in 2016, an amount greater than the total output of Texas or New York state. And the CBO’s forecast was, arguably, conservative—

⁵ In the presence of hysteresis, it is more plausible that shifts in demand could have long-run effects. The existence of a linear trend between 1947 and 2007 would then suggest that shifts in demand were not large enough or sustained enough to produce large effects on the economy’s productive potential.

In 2006, the CBO was forecasting cumulative 16 percent growth in real, per-capita GDP over the next decade. Instead, growth has been just 6 percent. If GDP had grown as predicted, an additional \$1.7 trillion would have been produced in 2016, an amount greater than the total output of Texas or New York state.

it was already predicting a significant slowdown in growth from what the U.S. had seen historically. Over the full postwar period, from 1947 to 2007, the U.S. saw an average 2.2 percent annual growth rate in real per capita GDP. If this rate had continued, we would have seen 25 percent growth over the past decade, and 2016 GDP would have been \$4.6 trillion higher than it actually was—additional goods and services equal to the output of New York, Texas, Illinois and Florida combined.

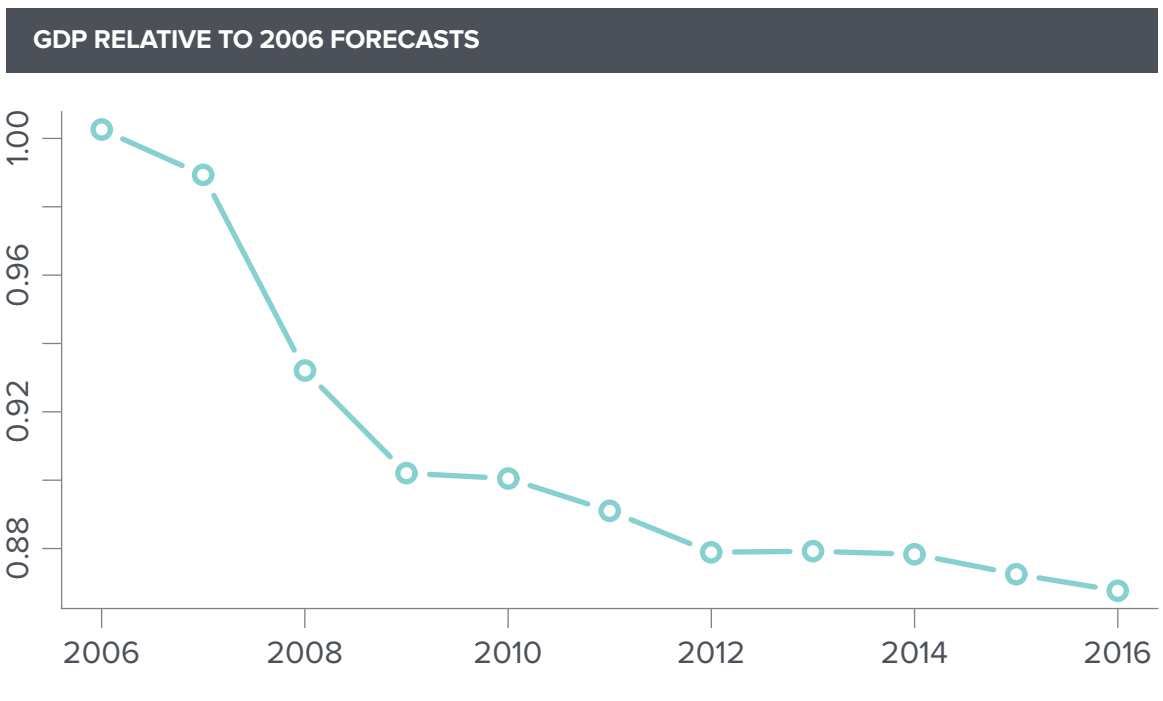


FIGURE 2 Real GDP as a fraction of the CBO's January 2006 estimate of potential GDP. The figure shows that the gap that opened up between actual GDP and forecast GDP during the recession, has not closed at all in the recovery. *Source: CBO, BEA, author's analysis.*

Figure 2 compares the path of real GDP to the CBO's estimates of potential GDP as of January 2006.⁶ Unsurprisingly, during the recession GDP fell well below the 2006 estimates of potential output—about 10 percent below by the official end of the recession. What is more striking, however, is that since then none of the gap with the 2006 estimate of potential has been closed. On the contrary, the gap continued to widen, so that by 2016 real GDP was approximately 12 percent below the CBO's 2006 estimate of potential output for that year.

Whether one takes the long-run trend or, more cautiously, the pre-recession forecasts as a baseline, this is a large, unanticipated and persistent loss in production. It implies substantially lower incomes for American households, and fewer resources available to address national challenges. It is certainly possible that this gap is entirely the result of overoptimistic growth forecasts prior to the recession. It may be, as many economists and policymakers seem to believe, that the post-recession slowdown in growth is a supply-side phenomenon, reflecting a diminished capacity to produce rather than any underutilization of existing capacities. It may be that this decline in potential has structural causes that are beyond the reach of policy, such as an aging population and an exhaustion of the technological frontier. (In particular, as we will see below, there is some empirical support for the argument that declining labor force participation reflects the changing age structure of the population—though the case is not as decisive as it is sometimes made out to be.)

By mistakenly regarding our underuse of existing labor, capital and technology as a shortage of real resources, we will miss out on the policy interventions that would allow us to use them more fully.

But there is danger in reaching these conclusions too quickly. Suppose that some or all of the shortfall is in fact due to a lack of demand—to not enough spending to fully use our existing capacities. In that case, there is a risk that macroeconomic policy itself will write into permanence what could have been a temporary interruption in the normal process of economic growth. By mistakenly regarding our underuse of existing labor, capital and technology as a shortage of real resources, we will miss out on the policy interventions that would allow us to use them more fully. These interventions are, conceptually if not practically, straightforward: someone, whether the government, business or households, needs to buy more.

⁶ The definition of GDP underwent a major revision in 2013, when intellectual property production and research and development spending were reclassified as investment spending. Because of this, the level of GDP cannot be directly compared with earlier forecasts. Instead, all comparisons of actual forecast output and productivity are of actual and predicted cumulative growth, using a consistent definition of GDP in each case. Since the revision did not affect employment data, employment levels can be compared directly.

1.1 U.S. GDP REMAINS AT ABOUT 15 PERCENT BELOW THE PRE-RECESSION TREND.

The first half of the 20th century—with the Great Depression and two world wars—saw very large swings in GDP growth. But after the demobilization following World War II, the U.S. has seen very smooth growth in output. Between 1947 and 2007, real GDP per capita grew at a steady 2.2 percent per year. As Figure 1 shows, this growth rate was remarkably consistent. If we define the output gap as the deviation between real GDP and the trend line, the standard deviation of the gap is 3.1 percent, and the median gap is just 2.5 points. The maximum gap in either direction was less than seven points, and those deviations were rare and short-lived. Over this sixty-year period, one could make very accurate forecasts of future GDP per capita simply by projecting the 2.2 percent annual growth forward. During the “Great Moderation” period from the mid-1980s to 2007, real output stayed even closer to trend, never departing from more than 5 percent in either direction and spending the majority of quarters within 1.5 points of it.

There is an important corollary to this steady long run growth rate: Periods of slower growth have alternated with periods of faster growth. Historically, every time a recession has reduced per-capita GDP relative to the long-run trend, it has been followed by a period of faster than usual growth. In other words, deviations from the trend have not been persistent.

Since 2007, however, there has been no tendency to make up the ground that was lost during the recession. On the contrary, the gap between per-capita GDP and the pre-recession trend has steadily increased, reaching 15 percent in the last quarter of 2016 (the most recent available). There is no precedent in the postwar period for a deviation from trend this large or this persistent.

Since 2007 there has been no tendency to make up the ground that was lost during the recession. On the contrary, the gap between per-capita GDP and the pre-recession trend has steadily increased, reaching 15 percent in the last quarter of 2016 (the most recent available). There is no precedent in the postwar period for a deviation from trend this large or this persistent.

OUTPUT GAP AND SUBSEQUENT GROWTH RATE

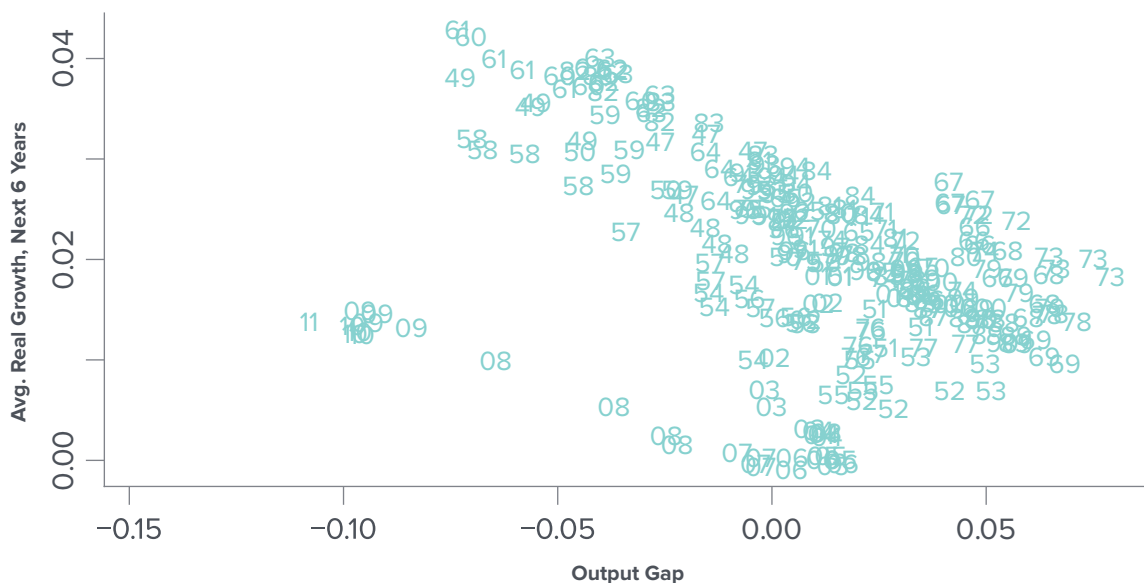


FIGURE 3 Output gap and real growth over next 6 years. The horizontal axis shows the output gap, measured as the deviation of real per capita GDP from the 1947-2007 trend line, in the year shown. The vertical axis shows the annual growth rate of real per capita GDP over the following six years. Data is quarterly. *Source: BEA, author's analysis.*

This is shown in Figure 3, which plots the output gap against the average real growth rate of per-capita GDP over the next six years. (Six years is about the average length of a U.S. business cycle.) The output gap is here defined as the deviation from the 1947-2007 trend. As the figure shows, prior to the most recent cycle, there was a very strong relationship between the output gap and the subsequent growth rate. The R-squared of the relationship is 0.6, meaning that the majority of medium-term growth could be predicted simply on the basis of the current deviation from trend. As the figure shows, every previous six-year period that began significantly below the long-run trend saw above-average growth, while almost every six-year period that began above trend saw below-average growth (the exception is periods beginning in the late 1960s and early 1970s).

This historical relationship would have led to predictions of unusually rapid growth in the following years, as growth returned to trend, as most forecasters in fact predicted. But as it turned out, not only was there no accelerated growth to bring output back to trend, the growth rates out of the recession did not even match historical averages. The only seven-year periods with comparably low per-capita GDP growth came in the wake of the booms around the Korean War in the 1950s, when unemployment fell to below 3 percent, and the late 1970s, periods during which output was well above trend and there were unambiguous

signs of overheating. There is no postwar precedent for an extended period of slow growth from a starting point that is already well below trend.

This deviation of GDP from its long-run trend is already, by itself, a problem for claims that the economy is now operating at full potential. There is no question that the financial crisis, collapse of the housing market, mass unemployment and general fear and uncertainty of the recession period itself depressed aggregate demand—that is, they reduced households' and businesses' desired spending. Because the sharp fall in output of 2008-2009 was, without question, driven by demand-side factors, logically there should have been some bounce back from the low levels of output at the bottom of the recession. The fact that we have not experienced the period of above-trend growth that characterized earlier recessions suggests that the demand problems have not been resolved.

There is no question that the financial crisis, collapse of the housing market, mass unemployment and general fear and uncertainty of the recession period itself depressed aggregate demand—that is, they reduced households' and businesses' desired spending ... The fact that we have not experienced the period of above-trend growth that characterized earlier recessions suggests that the demand problems have not been resolved.

1.2 MOST OF THE SLOWDOWN WAS NOT ANTICIPATED BY FORECASTERS. THE DECLINING OUTPUT GAP REFLECTS ESTIMATES OF POTENTIAL GDP CONVERGING TO ACTUAL GDP, RATHER THAN ACTUAL GDP RETURNING TO TREND.

A natural question to ask about this decline relative to the prior trend is whether forecasters anticipated it. This question is related to the underlying causes of the decline, since some factors, such as declining labor force participation due to an aging population, are easily anticipated by forecasters, while others, such as weak demand in the wake of a financial crisis, are not. For this reason, a number of authors use the fraction of the deviation from trend that was anticipated by prior forecasts as a measure of the contribution of structural

forces, and the fraction of the deviation that was not predicted as a measure of the contribution of weak demand or other effects of the financial crisis and its aftermath.

Ball (2014), Ball, DeLong and Summers (2014), and Fatas and Summers (2016), among others, try to answer the question of how much the decline in output is due to the recession by comparing pre-recession estimates of potential output with more recent ones for a number of rich countries. The three papers do different versions of this exercise, but all find that (1) the bulk of the slowdown in growth since the 2008-2009 recession was not predicted prior to it; and (2) there is no tendency for output to return to potential. Rather, changes in current output are eventually reflected in revised estimates of potential.

Despite the fact that the gap between per-capita GDP and the pre-recession trend has not closed, official measures of the output gap released by bodies like the Bureau of Economic Analysis (BEA) and CBO show a steadily closing gap between current output and potential. This is entirely because estimates of potential output have been steadily revised downward since 2007. In effect, the GDP ball has gotten closer to goal, but that is because the goalposts have been moving backward, not because the ball has been moving forward.

Estimates of potential output have been steadily revised downward since 2007. In effect, the GDP ball has gotten closer to goal, but that is because the goalposts have been moving backward, not because the ball has been moving forward.

The downward revisions result mechanically from the procedure used to estimate potential output. These estimates are based on a production function approach that includes “total factor productivity” (or the “Solow residual”), an unobservable parameter that acts as a kind of fudge factor. TFP growth is set at whatever level is required to make the estimate of potential output converge to actual output when unemployment is at its target level.⁷ The procedure followed by the statistical agencies cannot show a persistent output gap—the way the estimate is constructed ensures that actual output will always converge to potential over the business cycle. This is not necessarily a bad procedure for politically sensitive estimates whose function is to establish a consensus view for policymaking. But it’s important to understand that the decline in the official output gap provides no information about whether output is close to potential in any substantive sense.

⁷ For a thorough discussion of the CBO’s methodology for estimating potential output, see CBO (2001).

ESTIMATES OF POTENTIAL OUTPUT FOR VARIOUS DATES




FIGURE 4 CBO estimates/forecasts of potential output for various years. The heavy black line shows the historical path of GDP. The colored lines show historical and forecast estimates of potential output for each year from 2002 through 2016. The blue-green lines are estimates made before the recession, while the red-purple lines are estimates made after the recession. The 2006 value is set to 100. As the figure shows, the fact that GDP is now estimated to be close to potential is entirely the result of downward adjustments of potential estimates since the recession. *Source: CBO, author's analysis.*

The pattern of downward revisions is shown in Figure 4. The colored lines show the CBO's January estimates for potential output over the following 10 years.⁸ The heavy black line shows the actual path of GDP.

What do we see? First, there has been a systematic reduction in estimates of potential. While there are some upward adjustments in the early years, more recently all the adjustments have been downward. The estimates of 2015, for example, first made in 2005, have been reduced every year since then. The same goes for 2016 and all later years. These are not random errors, and they are not small. The estimate of 2016 potential GDP made by the CBO in 2006 was 10 percent greater than the estimate last year.

In 2006, before the recession, the CBO was predicting that real per-capita output would increase by 2 percent over the next decade—slightly below the long-run trend rate. In January 2008, the last pre-recession forecast, the CBO had revised its ten-year growth rate down slightly, to 1.8 percent. In the event, per-capita GDP grew at just 0.6 percent per year

⁸ Data is taken from the CBO's Potential GDP and Underlying Inputs for various years.

over the period—by far the slowest growth in per-capita output over any 10-year period since the 1930s. Of the 15-point gap between current per-capita output and the long-run trend, less than half was predicted before the recession. The majority of the gap reflects downward revisions since the recession began.

1.3 THE DECLINE IN GROWTH REFLECTS A MIX OF LOWER LABOR FORCE PARTICIPATION AND SLOWER PRODUCTIVITY GROWTH.

The next step in making sense of the GDP shortfall is to break it down into various components and see how they behave individually. There are a number of ways to do this, but the most straightforward is the identity:

$$(1) \text{ output} = \text{productivity} \times \text{employment} = \text{productivity} \times \text{labor force} \times (1 - \text{unemployment}) \times \text{hours per worker}$$

It follows that:

$$(2) \text{ Percent output growth} = \text{percent productivity growth} + \text{percent employment growth} \\ = \text{percent productivity growth} + \text{percent growth in non-institutionalized population} \\ + \text{change in the labor force participation rate} - \text{change in the unemployment rate} + \\ \text{percentage growth in average hours}$$

Output is employment multiplied by labor force productivity. This is also equal to labor force productivity multiplied by the labor force size multiplied by the employment rate (which is one minus unemployment rate). These are directly observable aggregates, and the CBO gives the relevant components for each of these in its estimates of potential output. Productivity here means labor productivity, or output per worker. As applied to potential, unemployment means the non-accelerating inflation rate of unemployment, or NAIRU. The unemployment rate is supposed to be consistent with stable inflation, which is targeted by the central bank.

The following table shows how these components have behaved over the past 10 years, compared with their long-run trends. It also shows the pre-recession forecasts of the CBO. The first column shows the difference between actual 2016 output and what you would predict by projecting forward the 1989-2006 trend.⁹ The second column shows the CBO's

⁹ The trend starts in 1989 because that year, like 2006, had actual output very close to estimated potential. Starting there, rather than in an earlier year, means that we can ignore population growth, since population increased at almost exactly the same rate between 2007 and 2017 as between 1989 and 2007.

2006 forecasts for 2016. The third column shows the value you would predict for 2016 by simply extrapolating the 1989-2006 trend. The fourth column shows current estimates of potential for 2016. And the final column shows actual values for 2016. In the first four rows, the 2006 value is set to 100. (Because actual output in that year was very close to potential, it doesn't matter whether we use actual or potential output for 2006.)

2006-2016 GDP GROWTH: TREND, FORECAST AND ACTUAL					
	Actual 2006	2006 Forecast for 2016 Potential	Trend 2016	Revised 2016 Potential Estimate (CBO Estimate)	Actual 2016
GDP per capita	100	116	125	107	106
GDP	100	126	135	116	114
Employment	100	107	115	–	106
Productivity	100	118	118	–	108
Labor force participation	66.2	62.9*	66	–	62.8
Unemployment rate	4.4	5.2	–	4.8	4.8

TABLE 1 GDP and components, 2006 and 2016, forecast and actual. “Trend” is an extrapolation of the 1989-2007 growth rate. “Forecast” is the value given or implied in the CBO’s 2006 “Potential GDP and Underlying Inputs.” Potential unemployment is the CBO’s current NAIRU estimate. *Source: CBO, BEA, author’s analysis.*

***Note:** The CBO’s annual forecasts do not directly report a labor force participation rate; their forecast includes only potential employment. CBO (2001) indicates that CBO computes potential employment based on the Social Security Administration’s forecasts of the working-age population. So the “forecast” labor force participation rate reported here combines the CBO’s forecast of employment with the 2016 labor force forecast in the 2007 OASDI Trustees Report. However, a separate 2007 report does include a forecast for the labor force participation rate of 64.3 percent for 2016, significantly above the actual figure (CBO 2007), so it is possible that the accurate forecast of employment growth in fact reflects an overestimate of labor force participation combined with an underestimate of population growth.

As the table shows, GDP in 2016 was about 15 percent below what we would predict based on the earlier trend. Prior to the recession, the CBO predicted 7 points of this slowdown, or a bit less than half. But the picture is quite different for employment and productivity taken separately.

The deceleration in labor force growth was almost entirely predicted by the CBO. The productivity slowdown was not predicted at all. The downward revisions to estimates of potential output over the past decade have almost all been revisions to productivity estimates. (There were important downward revisions to forecasted labor force growth, as shown in Figure 5, but they came a bit earlier.) Unemployment does not play an important role in either case: both actual unemployment and the current estimate of NAIRU are very close to the 2006 unemployment rate. This is quite different in Europe, where higher NAIRUs explain a large part of the change in potential output (Ball 2014).

ESTIMATES OF POTENTIAL EMPLOYMENT FOR VARIOUS DATES




FIGURE 5 CBO estimates/forecasts of potential employment for various years. The heavy black line shows the historical path of employment. The colored lines show historical and forecast estimates of potential output for each year from 2002 through 2016. The blue-green lines are estimates made before the recession, while the red-purple lines are estimates made after the recession. The 2006 value is set to 100. As the figure shows, cumulative employment growth since 2006 is now close to the CBO's pre-recession forecasts. *Source: CBO, author's analysis.*

ESTIMATES OF LABOR PRODUCTIVITY FOR VARIOUS DATES




FIGURE 6 CBO estimates/forecasts of potential labor productivity for various years. The heavy black line shows the historical path of productivity. The colored lines show historical and forecast estimates of potential output for each year from 2002 through 2016. The blue-green lines are estimates made before the recession, while the red-purple lines are estimates made after the recession. The 2006 value is set to 100. *Source: CBO, author's analysis.*

ESTIMATES OF NAIRU (UNEMPLOYMENT) FOR VARIOUS DATES




FIGURE 7 CBO estimates of the non-accelerating inflation rate of unemployment (NAIRU) and the actual unemployment rate, various dates. *Source: CBO, author's analysis.*

These different patterns of revisions are shown in Figures 5 to 7. The format is the same as for Figure 3—each line represents the estimates made in a given year, with the bluer lines showing earlier forecasts and the redder lines later ones.

Figure 6 shows actual and forecast labor productivity, defined here as output per worker. The figure shows there are modest downward adjustments in 2007 and 2008 but the big adjustment comes later, especially in 2009 and 2014. And the later adjustments are not just to the level of productivity but also to the trend. So not only is there no convergence between actual productivity and pre-recession forecasts, the gap has continued to widen over time. For the labor force, by contrast, the biggest adjustments come before the recession, especially in 2003, when the trend is revised downward. The post-recession revisions are small, and the actual trajectory of the labor force shows a definite reversion toward the immediate pre-recession forecasts.

Labor productivity can be further subdivided into output per hour and hours per worker. We focus on output per worker rather than output per hour in Table 1 and the related analysis, for two reasons. First, the CBO forecasts, which are our baseline, are constructed in terms of output per worker, with no separate hour term. Second, reduced work hours play no role in the post-2007 deceleration in labor productivity. While hours per worker did fall somewhat

during the recession, they quickly recovered. Both weekly and annual hours per worker are now approximately the same as they were in 2007, a level that has been roughly constant since the early 1980s. Consequently, the results reported above would be essentially the same if productivity were calculated per hour rather than per employee.

Output per hour does diverge from output per worker in two important ways, however. First, while hours per worker have shown no trend since 1980, during the preceding three decades, average annual hours declined at around a third of a percent per year. This slow but steady decline translated into an 11 percent decline in annual hours per worker between 1951 and 1980. More importantly for present purposes, it means that output per worker understates the difference in labor productivity between the postwar decades and the recent period. In effect, the decline in hours in the 1950s, 1960s and 1970s created a headwind for growth that is no longer present today. Second, while there is no overall trend in hours worked between 1980 and today, there is a substantial rise—from around 1800 hours per year up to 1,850—during the second half of the 1990s, which was subsequently reversed. At around 0.5 points per year, this increase in hours accounts for about half the anomalous acceleration in labor productivity during this period. If, as seems likely, this increase in hours was mainly a response to stronger labor market conditions, this is another important reason for seeing the late 1990s boom as in large part a demand-side phenomenon. If strong demand elicited a significant increase in hours per worker then, that suggests hours may be a margin on which continued expansionary policy could boost output today.

If strong demand elicited a significant increase in hours per worker then, that suggests hours may be a margin on which continued expansionary policy could boost output today.

Finally, the estimated NAIRU was adjusted upward during the recession and back down since then with no sustained movement one way or the other. It's worth noting, however, that while there has been no systematic adjustment of the NAIRU, there have been substantial short-term changes. From the late 1990s until 2007, the CBO based its potential output estimates on a fixed NAIRU of 5.2 percent—an unemployment rate above this implied output below potential, while an unemployment rate below this implied output above potential. In 2007, the CBO changed its NAIRU estimate to 5 percent, and in 2008 to 4.8 percent. These new NAIRU estimates were applied retroactively, resulting in an upward adjustment to estimated potential output for earlier years. Then, starting in 2010, the CBO began adjusting its NAIRU estimate annually, based in part on current unemployment rates. Between 2009 and 2014,

the NAIRU estimate was adjusted upward by a full 1.2 points, from 4.8 to 6.0, before being adjusted back downward. This mix of upward and downward adjustments does not play an important part in the long-run shifts in estimates of potential. But it is an object lesson of the dangers of being too quick to interpret a demand shock as a shift in the economy's productive potential. In retrospect, it is clear that the rise in unemployment starting in 2008 did not include any important structural component; unemployment today is back to its pre-recession levels, with no sign of inflation. But if the CBO's estimates had been the basis for policy, the fall in unemployment would have been halted sometime in 2007, and nearly two million people would have been condemned to joblessness instead of finding productive work.

Returning to the continuing gap between post-recession output and the pre-recession trend, we see it combines two distinct developments: (1) a widening productivity gap, which was not predicted by the CBO, and (2) a narrowing labor force gap, which was mostly predicted.

Figure 8 shows five-year rolling averages of output, employment and labor productivity growth since 1947. As of 2016, slower productivity growth and slower employment growth contribute about equally to the shortfall of output from the earlier trend. If we want to figure out the relative role of the demand and supply sides in the shortfall, we need to look at these two components individually.

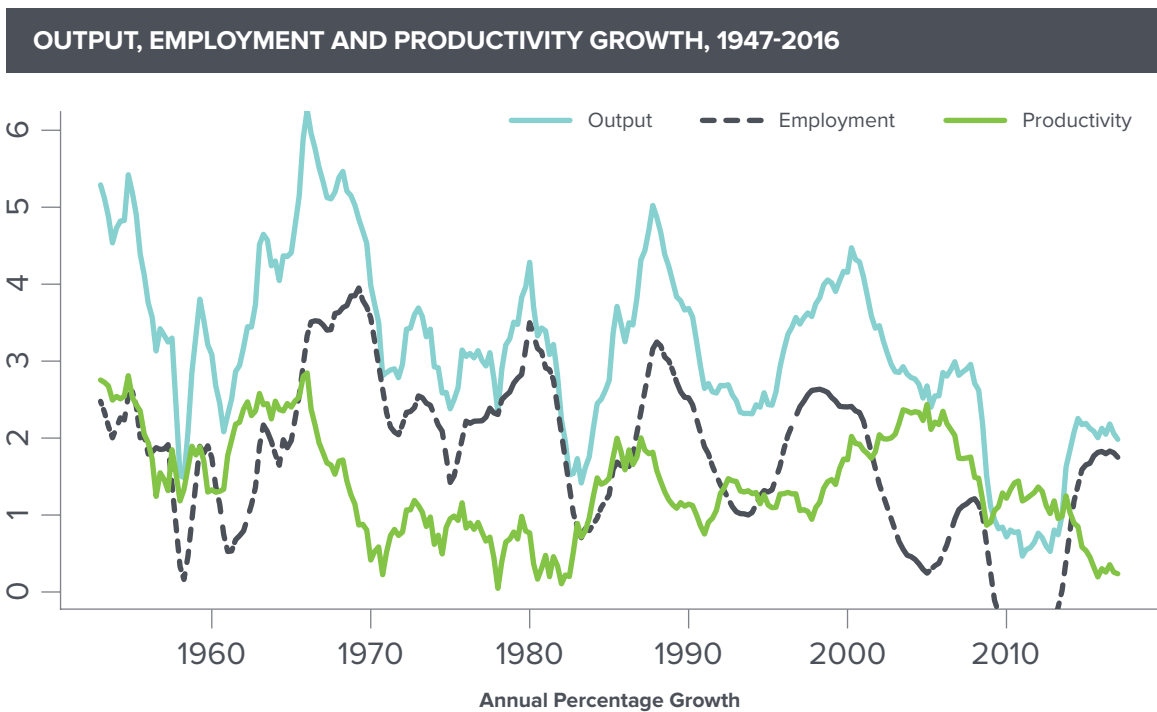


FIGURE 8 Output, employment and productivity growth, rolling five year averages. This shows average annual growth in real GDP and the two components for five-year periods ending in the year shown. *Source BEA, author's analysis.*

HOW MUCH WEIGHT SHOULD GROWTH FORECASTS GIVE TO THE 1990s?

It is sometimes argued that the shortfall of productivity growth relative to the CBO's pre-2007 forecasts is simply due to earlier forecasts mistakenly projecting the exceptional growth of the "tech boom" period of the late 1990s into the indefinite future. It is true that 1997-2003 saw exceptionally rapid productivity growth, averaging 2.3 percent a year, compared with a postwar average of 1.5 percent (See Figure 8). Given the essentially extrapolative nature of CBO's forecasts, this rapid growth inevitably influenced its long-run forecasts, which as of 2007 saw annual productivity growth converging to 2 percent. While lower than the peak rates reached in the 1990s, this is still faster than in most other periods since the 1960s. The claim that the CBO's forecasts were influenced by the exceptional productivity performance of the late 1990s certainly has an element of truth.

But this is not a sufficient explanation for the shortfall of productivity relative to forecasts, for three reasons. First, the CBO's pre-recession forecasts already explicitly took note of the exceptional character of this period. Forecasts made in the mid-to-late 2000s already incorporated a special ad hoc "temporary faster growth" term that added about a quarter of a point to annual productivity growth for the late 1990s and early 2000s, that was not incorporated into

forecasts of future growth. Second, the 1990s effect can explain only part of the shortfall. Since the end of the recession, labor productivity growth has averaged 0.7 percent per year, compared with the CBO's forecast of 2.1 percent. If we look at the 1980-2007 period, but exclude the years 1998 through 2002 (the period of rapid productivity growth), average labor productivity growth is still around 1.5 percent a year. Even if the CBO had entirely ignored the tech boom period, 60 percent of the productivity shortfall (0.8 out of 1.4 points) would still be present. (See Figure 8.)

Third, and perhaps most important, the CBO was right to take the tech boom period into account in forecasting future growth. The experience of this period is informative about the growth potential of the U.S. economy. Even if we recognize that this period was in some ways exceptional, it makes no sense to put zero weight on it in forming our forecasts of the future. At minimum, the years 1998 to 2002 should be given the same weight as other years in estimates of potential productivity growth. If the question is the potential of the economy, the years of best performance are arguably more relevant than others.

This argument grows stronger when we recall similar debates 20 years ago.

By 1997, many policymakers believed that the economy was at potential, and that the Fed needed to tighten to avoid dangerous overheating. Alan Greenspan made the politically risky choice to resist this pressure, and to maintain low rates in the belief that more rapid growth was possible. The Fed's willingness to maintain low rates even as unemployment fell below 4 percent—opposed at the time by much of the FOMC—is one of the key events in recent macroeconomic history. It is only because of this decision that the low-unemployment, high-productivity economy of the later 1990s came into being. (The ways in which stronger demand can contribute to higher productivity are discussed in Section 3.1 below.) Far from being a one-off event, this has obvious relevance to the choices facing the Fed today. This period—the one that we are told we should ignore—is precisely the last period in which the Fed was facing a choice similar to the one it confronts today.

So while it is true that the productivity shortfall looks bigger relative to a trend that includes the 1990s, this is a strange argument to make against continued expansionary policy today. In effect, it argues that when we consider the potential costs and benefits of a high-pressure economy today, we should not take into account the evidence from the last time the experiment was actually tried.

Labor force participation

Between 2006 and 2013, the headline labor force participation rate fell from 66 to 63 percent, where it has since remained. Since unemployment had returned by 2016 to its 2006 levels, it is this fall in labor force participation that drove the decline in the employment-population ratio described in the previous section.¹⁰ A central question about this decline is the extent to which it was the result of the aging of the population, or of demographic changes more broadly. If demographic factors were dominant, then the smaller share of the population in employment is not a sign of weak demand. Rather, it is the predictable result of the tendency of older people, by choice or necessity, to exit the workforce. This story is consistent with the fact that most of the decline in labor force participation was already predicted by the CBO a decade ago. But the fact that the decline in labor force participation was predicted does not mean that it was inevitable, or independent of demand conditions or policy more broadly. Before assuming that the decline in employment is mostly or entirely a question of demographics, we need to look at the data more carefully.

Before assuming that the decline in employment is mostly or entirely a question of demographics, we need to look at the data more carefully.

It is certainly true that the segment of the population over 65 has been growing rapidly over the past decade. It is also true, though less widely recognized, that labor force participation rates historically decline steeply as people pass through their 50s, and the share of the population in the upper part of this range has also increased. This raises two questions: First, how much of a decline in labor force participation could we have predicted on the basis of demographic changes in the population, all else equal? And second, how hard a constraint on employment is this demographic change?

To answer the first question, we calculate employment-population ratios for different demographic groups in a given year. Then we apply those same rates to the different shares of the demographic groups in subsequent years. The result tells us what the change in employment would have been, given the actual change in the composition of the population,

¹⁰ For a given population, the employment-population ratio is equal to the labor force participation rate times one minus the unemployment rate.

if participation rates within each group had remained constant. We can think of this as the change in labor force participation “explained” or “predicted” by demographics. The remaining, within-group component is the change in employment rates by people within given demographic groups; it cannot be explained by demographic changes, but must be due to something else. This approach, widely used to analyze changes in a heterogeneous population, is known as a shift-share analysis.

While the basic approach is uncontroversial, we face a couple of questions before we can put it into practice. First, which demographic categories do we consider? And second, what year should we use as a base year? Fortunately, with annual data, we don’t need to pick one year. Instead, we can “chain” together different years, in the same way that national statistical agencies do when constructing inflation measures. For each year, we can calculate the part of the change in employment that is explained, in a statistical sense, by changes in demographics, and the part of the change that is within individual demographic groups. Then we string together these changes to generate two counterfactuals. One shows what would have happened to employment if the demographic mix of the population had changed the same as it actually did, while employment rates within each group remained constant. The other shows what would have happened if employment rates within each group had changed as they actually did, while the demographics remained constant. (There is also a residual term. But it is normally quite small, so the fixed-rates-within-group and fixed-composition changes add up to close to the actual change.)

This chaining approach is followed by other recent studies of changes in labor force participation rates (Aaronson et al. 2014). It has the advantage that there’s no need to make an arbitrary choice of any particular base year. Unfortunately, there is no such easy fix to the question of which demographic groups to include.

Age has a large effect on employment status, and the age profile of the U.S. population is shifting significantly over time. But are there other demographic factors we should also consider?

Most discussions of demographic sources of the decline in the employment-population ratio focus on age. This makes sense: Age has a large effect on employment status, and the age profile of the U.S. population is shifting significantly over time. But are there other demographic factors we should also consider?

Employment status also varies strongly by sex. But unlike age, the sex ratio of the population (at least in the U.S.) does not vary significantly over time. Another important but less discussed demographic characteristic relevant to employment is the presence of children in the household. Mothers of young children are historically less likely to work, while fathers are more so. The first factor has significantly outweighed the second, and so a larger number of families with young children implies a lower overall employment rate. The number of families with young children has declined over time, especially in the immediate aftermath of the recession.

Another important but less discussed demographic characteristic relevant to employment is the presence of children in the household. Mothers of young children are historically less likely to work, while fathers are more so.

Finally, we might consider education. People with high school and college degrees are consistently more likely to be employed than people without, independent of variables such as age, sex and family situation. But it is not obvious that we should consider this relationship between education and employment as “demographic” in the same way we consider age or the presence of children. Race is also often included in attempts to distinguish the demographic component of changes within the employment-population ratio. It is true that employment status varies by race, with African-Americans, in particular, significantly less likely to be employed. But given the overwhelming evidence that racial discrimination remains important in labor markets, it seems inappropriate to include race as a demographic variable. Unlike age, racial variations in employment rates do not reflect variation in capacity or willingness to work. Rather, they reflect pervasive queuing in labor markets, in which people of color are the last hired and first fired.

In our first analysis, we include only age and sex as demographic variables. (The issue of children is discussed at the end of the section.) The variables interact, so we have an employment rate for each age-sex combination: men aged 34, for example. Because there are good arguments both for including and excluding education as a demographic variable, we calculated both sets of counterfactuals—one with it, and one without. The results are shown in Figures 9 and 10—the first without education, the second including it.

The conclusion of both analyses is the same: Whether you start from the all-time high point in 2000 or the most recent business-cycle peak in 2007, the majority of the decline in the employment-population ratio is due to lower participation within demographic groups, not changes in the composition of the population.

DECOMPOSED CUMULATIVE ANNUAL CHANGES IN EMPLOYMENT-POPULATION RATIO



FIGURE 9 Decomposition of cumulative change in the employment-population ratio (employment rate) into demographic shifts and within-group changes. The blue line shows the change in the employment rate that one would predict based on changes in the age and sex composition of the population, given constant employment rates within age-sex groups. The red line shows the change in the employment rate that would have occurred if the age-sex composition of the population had remained fixed. The green line shows the actual change in the employment rate; it is approximately equal to the sum of the other two lines. *Source: CEPR, Current Population Survey, author's analysis.*

Figure 9 shows the two components using age and sex only. What do we see?

Since 2000, the U.S. employment-population ratio has declined by about five points, of which four points has come since 2007. Out of that, three points (2.5 points since 2007) have reflected lower participation rates within demographic groups, while two points (1.5 since 2007) reflect changes in the composition of the population. Measured this way, at most 40 percent of the decline in labor force participation reflects demographic changes, while 60 percent must be due to other factors. The conclusion that at most half the decline in labor force participation is attributable to aging is similar to that reached by the Council of Economic Advisors when it explored the same question (CEA 2014).¹¹

¹¹ The somewhat higher demographic share reported by the CEA is a result of their using age and sex and race, rather than age, sex and the presence of young children, as explanatory variables. When we add race as a demographic variable, the demographic share of the decline in the employment ratio rises to about 60 percent, very close to the value reported by the CEA. But we do not think that it is appropriate to use race in this way, for reasons discussed in the text.

CHANGE IN EMPLOYMENT-POPULATION RATIO (MARCH 2000 – MARCH 2016, BY DEMOGRAPHIC GROUP)

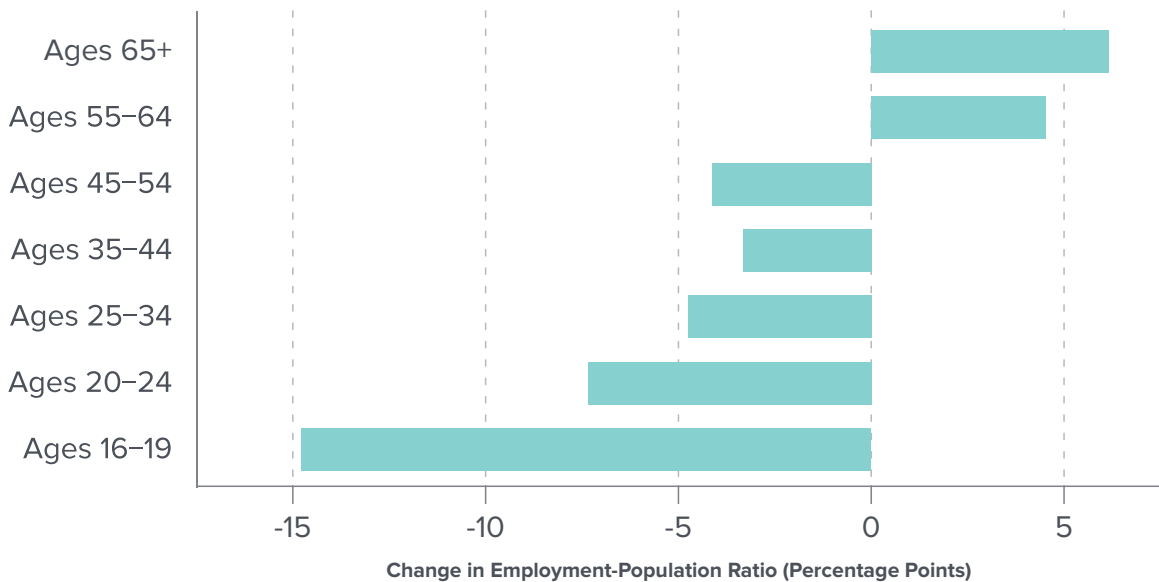


FIGURE 10 Change in employment population ratio by age group, March 2000 to March 2016. Source: CEPR, Current Population Survey, author's analysis.

To get a clearer picture of the within-group component of the fall in the employment-population ratio, Figure 10 shows the change in the employment-population ratio for each age group.¹² As the figure shows, the within-group component varies strongly by age, with the largest falls in employment ratios among younger groups and employment rates among the oldest groups actually increasing. The fall among the younger groups may reflect a higher proportion of higher education, while the rise among the older groups may reflect the decline of private pensions, as well as the increase in the normal retirement age under Social Security. However, the pattern is also consistent with a story of continued labor-market weakness. Young people with limited work experience are likely to be at the back of the line for employment, and their choices about participation in the labor market are most likely to be sensitive to demand conditions. Older people, on the other hand, may respond to weak labor markets and economic insecurity by putting off retirement, especially if their incomes have fallen or if spouses or other family members are unable to find work.

We can't explore further the relative contributions of these different factors. But the evidence here does strongly suggest that, with respect to employment, demographics are not destiny. Even to the extent that the fall in employment rates is apparently accounted for

¹² See the similar discussion by Doug Henwood at LBO-News: <https://lbo-news.com/2017/04/18/how-employed-are-we/>

by demographic changes, this does not reflect a fixed demographic pattern of employment applied to a shifting population, but rather offsetting change within different groups. Given the large shifts in employment rates within demographic groups, it is hard to believe that labor force participation would not respond to stronger demand.

What if we add education as a demographic variable? This is common in labor economics and there is a natural logic to it. Historically, people with more years of schooling have been consistently more likely to be in the labor force than otherwise similar people with fewer years of schooling. Economists often think of education as a kind of “human capital” that allows workers to be more productive and increases their earnings. It’s natural to extend that to see education as also functioning to overcome barriers that might prevent people from engaging in paid work at all. People with more education may also feel a stronger desire to engage in paid work. Consequently, it’s reasonable to suppose that, just as we would expect an older workforce to have fewer people engaged in paid labor, all else equal, a more educated workforce should have more people engaged in paid labor.¹³

The results of adding education as a demographic variable are shown in Figure 11.

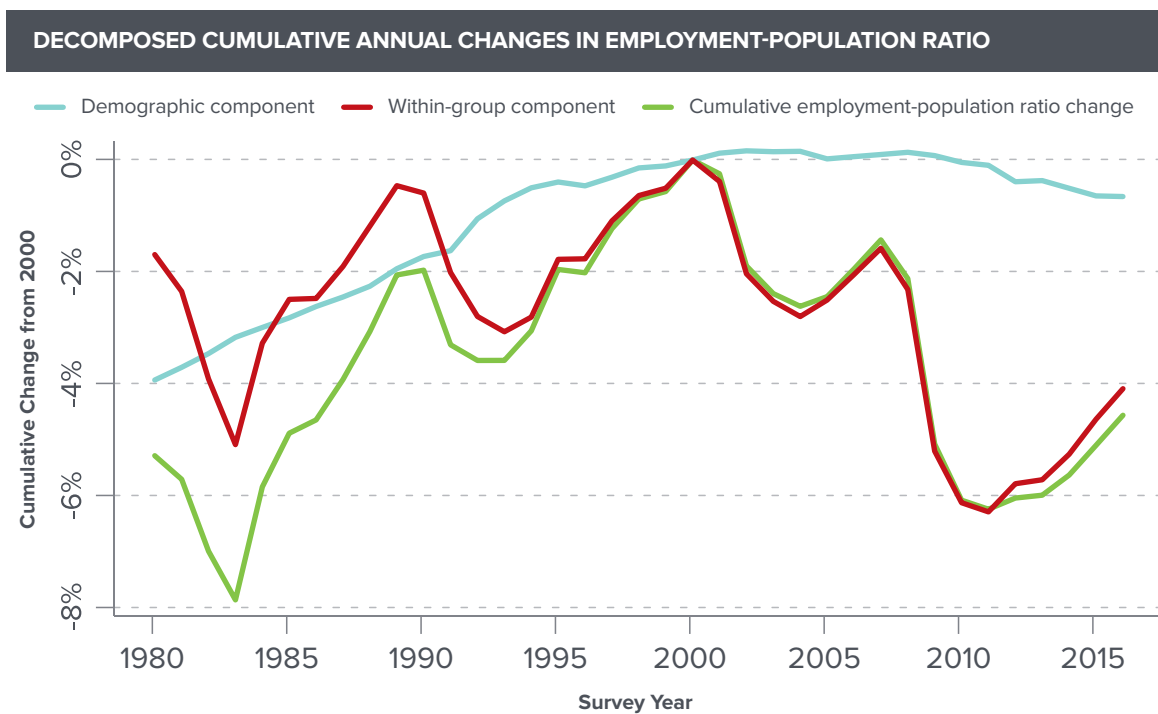


FIGURE 11 Decomposition of cumulative change in the employment-population ratio (employment rate) into demographic shifts and within-group changes. The blue line shows the change in the employment rate that one would predict based on changes in the age, sex, and education composition of the population, given constant employment rates within age-sex-education groups. The red line shows the change in the employment rate that would have occurred if the age-sex-education composition of the population had remained fixed. The green line shows the actual change in the employment rate; it is approximately equal to the sum of the other two lines. *Source: CEPR, Current Population Survey, author’s analysis.*

¹³ Alternatively, or in addition, education can be included in estimates of potential output as a measure of “labor quality”. (Hall 2016) In that case, rising education makes the productivity slowdown even more puzzling.

As Figure 11 shows, once we include education, the demographic component essentially vanishes, accounting for less than half a point of the five-point decline in labor force participation. If we define demographic groups by age, sex, and education, over 90 percent of the decline in participation has been within demographic groups. The changing composition of the population has had a negligible effect.

In other words, while it is true that the American population is getting older, it is also growing more educated. All else equal, older people are less likely to be employed, and educated people are more likely to be. Over the past decade or two, these two factors have essentially canceled one another out.

It is true that the American population is getting older, it is also growing more educated. All else equal, older people are less likely to be employed, and educated people are more likely to be. Over the past decade or two, these two factors have essentially canceled one another out.

Which of these results is more meaningful? The statistical argument for including education is the same as for including age, sex or presence of children—all of these characteristics have historically been associated with significant differences in employment rates. And there's nothing controversial about saying that increased education makes people both more willing and able to find work. So perhaps the second graph, with its negligible demographic contribution, is the one we should focus on. There is a strong sense that people's employment status follows a "natural" progression over a lifecycle first a period of education, then a period of labor, and finally a period of retirement. Some people similarly believe that caring for young children naturally implies spending a substantial period outside the labor force, though this belief is not as prevalent as it once was. Both aging and child rearing are deeply rooted aspects of life that one would not expect to respond quickly, if at all, to changes in labor market conditions. But no one sees the link between education and employment in these quasi-biological terms. And while it is certainly true that less educated people are less likely to be employed, it is not clear how much of this is coming from the supply side. It may be that people with weaker educational credentials are willing and able to do the same kinds of jobs but find themselves at the back of the line.

It is worth noting, that even if changes in employment rates do in some sense reflect changes in demographics, that does not mean they are inevitable or unrelated to demand or policy variables. We should not assume there is a fixed, "normal" participation rate for a

given demographic group. The large increase in women’s labor force participation between the 1950s and 1980s, for instance, would not have been predictable on demographic grounds. And while the postwar U.S. has not seen sustained demand shocks on the level of the Great Recession, history does suggest that large demand shifts can have large, persistent effects on labor force participation rates. The fall in employment in the 1930s, for instance, reflected not only elevated unemployment rates but also a large fall in measured participation rates; participation then rose even more sharply with the “effectively unlimited labor demand” during the war (Mathy 2015). Arguments implying that the U.S. today is at potential would, in 1940, have ruled out the large increase in output and employment that in fact took place during the war.

We find that changes in the composition of the population can explain as much as 40 percent of the fall in the employment-population ratio since 2007, if the effects of rising education are ignored. If demographic variables are defined to include education, the contribution of demographic change to falling employment is essentially nil.

In conclusion, we find that changes in the composition of the population can explain as much as 40 percent of the fall in the employment-population ratio since 2007, if the effects of rising education are ignored. If demographic variables are defined to include education, the contribution of demographic change to falling employment is essentially nil. In any case, the majority of the fall in employment rates to date has taken place within demographic groups—it would have occurred even if the composition of the population were fixed. And even the upper limit of 40 percent must be treated with caution. Historically, demographic variations in labor force participation have not been a hard constraint, but have been responsive to a range of social and policy factors, including the state of the labor market.

If the fall in labor force participation is due, at least in part, to the lowering effects of the recession, or to weak demand more generally, that means there is more space for expansionary policy than conventional measures suggest. It also means the social benefit of such policies is greater. As Ball (2015) suggests, a “high-pressure economy” will not only boost output and employment in the short-run, it will disproportionately benefit the chronically unemployed, and will strengthen enduring ties to the labor market for millions of people who were previously excluded from it.

DOES THE NUMBER OF FAMILIES WITH CHILDREN MATTER FOR EMPLOYMENT?

EFFECT OF CHILDREN ON EMPLOYMENT PROBABILITY




FIGURE 12 Effect of presence of young children on the probability of employment various years. The figure shows the regression coefficient of the presence of young children in the household on the employment of adults. In the early 1990s, parents of young children were about 10 percent less likely to be in the labor force; but by 2010, this effect had largely disappeared. *Source: NBER, Current Population Survey, author's analysis.*

Because of data limitations, we did not include presence of children in our main analysis. But as it turns out, adding a variable for young children does not change the results.¹⁴ While including education as a demographic variable has dramatic effects, the effect of including children is much smaller. On the face of it, this is surprising. There has been a sharp drop in the proportion of families with young children over the past decade. Based on historical evidence, one would

expect this to lead to a significant increase in women's participation in the labor force.

Why don't we see this effect in the data? While there has been a sharp drop in the percentage of U.S. households with children under the age of six over the past decade, there has also been a long-

variable for presence of own children under the age of six. The results were essentially unchanged.

run convergence between employment of parents and non-parents, as shown in Figure 10. In the early 1990s, the coefficient on the presence of children under six was, controlling for the other demographic variables, around -0.1. Parents of young children were about 10 percent less likely to be employed than demographically similar non-parents. But in recent years, the coefficient is around -0.01, meaning that the presence of a young child reduces the estimated employment rate by only one percent. So the decline in the share of families with children offsets only a trivial part of the decline in participation attributable to age.

Over the past generation, a shift in social norms around parents, especially mothers, engaging in paid work ... has diminished the negative effect of motherhood on employment to the point that it is almost entirely offset by the positive effect of fatherhood.

This is important, first, to have a clear picture of what demographic factors are relevant in establishing a baseline change in employment rates. Second, it is a demonstration of how “demographic” factors are conditioned by social and

policy changes. Over the past generation, a shift in social norms around parents, especially mothers, engaging in paid work plus the country’s fitful, inadequate, but still substantial progress toward provision of childcare has diminished the negative effect of motherhood on employment to the point that it is almost entirely offset by the positive effect of fatherhood. (The strong labor market of the late 1990s may well have contributed to cementing this shift, though this is not obvious in the data.) It is easy to imagine a similar reduction in the tendency of adults in their 50s and 60s to exit the labor force, given appropriate social supports (especially for the partially disabled) and a strong labor market.

Productivity and demand

There is good reason to believe that the slowdown in labor productivity is largely or entirely a reflection of weak demand.

Economics textbooks often treat productivity as “exogenous”, or fixed by forces outside the economy, such as technology or natural resources. While these sources of productivity grow over time, they don’t respond directly to changes in aggregate demand. A decision by government, business or households to spend more does not, in this story, do anything to improve the capabilities of businesses to produce more. More spending can of course mobilize resources that are currently unused, in particular unemployed labor. But once the unemployed workers have found jobs, in this view, increased spending cannot call forth and further additional output.

In the real world, things are not so simple. There is not always a hard line between utilizing previously unused resources and improving productive capabilities. Workers can switch from less to more productive tasks or occupations, or work more intensely; capital and other resources can similarly be used more intensively, or reallocated to higher value activities. Higher wages and other costs may encourage businesses to abandon low-productivity products and processes, or to develop or adopt higher-productivity ones. Technological advances do not come from the sky, but from the concrete activities of innovation and entrepreneurship that are more easily carried out during a boom than in a slump. In this section, we discuss a number of reasons to expect stronger demand to be associated with faster productivity growth. And we present some evidence that the unprecedented productivity slowdown of the past decade may be at least partly due to the period’s exceptionally weak demand.

3.1 THERE ARE MANY REASONS TO THINK THAT DEMAND COULD AFFECT PRODUCTIVITY.

Economics textbooks normally treat labor productivity as determined by technology and real resource endowments—things that don’t change simply because someone spends more money. Indeed, in a textbook world, we should expect productivity to fall when demand rises, and vice versa. Increased spending can quickly draw workers out of unemployment, but it cannot increase the stock of capital goods and other resources so quickly, or at all. And, in the textbook world, the most skilled workers are presumably the first to be hired, with the

pool of the unemployed representing the less qualified. So in a boom, the marginal worker is less skilled, and has less capital to use, meaning he or she should produce less. At the bottom of a slump, when unemployment is high, the least-skilled workers have been laid off and the remainder has more capital to work with. Consequently, productivity should be higher.

Labor productivity almost always falls in recessions, and rises in the recovery.

It is well known that business cycles do not look like this. As discussed in 3.3 below, labor productivity almost always falls in recessions, and rises in the recovery. This is presumably due to the costs of hiring and firing. Changes in spending do not immediately cause workers to move in and out of employment. It also reflects the fact that real-world production processes involve specialized equipment and facilities that can't easily be spread across a smaller number of workers. Capital as well as labor is idle in a downturn.

These factors, however, are usually understood to be strictly short-term, relevant over a year or two at most. Are there any plausible stories about how weak demand might depress productivity growth over a decade?

In fact, there are a number of reasons why productivity may respond to demand both in the short run and over extended periods. While this idea is not normally part of macroeconomic policy discussions, it is common sense in many other settings. (For instance, both supporters and opponents of minimum wage increases share the idea that higher wages may raise labor productivity.) It is likely that some of the fall in productivity growth would have happened independently of the Great Recession and the subsequent period of weak demand. A disproportionate share of the productivity slowdown is accounted for by the electronics, telecommunications and information sectors. This may well reflect technological (or measurement) issues specific to these industries rather than macroeconomic conditions. But, as discussed in Section 3.2 below, the productivity slowdown is widespread; most of it is accounted for by industries that did not see any acceleration of productivity in the 1990s. There are quite a few plausible ways this broadly distributed productivity slowdown could be explained by weak demand. Here are several stories to explain this.

The productivity slowdown is widespread; most of it is accounted for by industries that did not see any acceleration of productivity in the 1990s.

POSSIBLE LINKS FROM DEMAND TO PRODUCTIVITY

Increased hours	If we measure productivity as output per worker, as we usually do, then an increase in average hours worked will show up as an increase in productivity. There is a cyclical component to this. In recessions, employers reduce hours and lay off workers. According to the Bureau of Labor Statistics (BLS), seasonally adjusted weekly hours fell from 34.4 prior to the recession to a low of 33.7 in the summer of 2009. While a 2 percent fall in hours might seem small, it's a big change in less than two years, especially when you consider that real output per worker normally rises by less than 2 percent a year. But over longer periods, and in particular for the post-2007 period we're mainly interested in, it has not played a significant role. Changing hours seem to have been quite important in the 1930s and 1940s, but since then they have explained only a small fraction of variation in output per worker.
Less pseudo employment	In any economy, there are activities that are formally classified as jobs but are not employment in any substantive sense. You can take these "jobs" without anyone making a decision to hire you, and they don't come with a wage or any similar claim on any established production process. Joan Robinson was one of the first economists to discuss this; she called these kinds of pseudo-jobs "disguised unemployment." Her examples were someone who gathers firewood in a poor country, or sells pencils on street corners in a richer one. You could add work in family businesses, including farms, and various kinds of self-employment and commission-based work to this category. These activities absorb people who are unable to find formal jobs; the marginal product of additional workers here is normally very low. So if higher demand draws people from this kind of disguised unemployment back into regular jobs, measured productivity will rise. ¹⁵
Increased work effort	Because hiring and firing is costly, business don't immediately adjust staffing in response to changes in sales. When demand falls, businesses will initially keep some redundant workers because paying them is cheaper than laying them off and replacing them later; and when demand rises, businesses will first try to get more work out of existing employees rather than paying the costs of hiring more. Some of this takes the form of the hours adjustment, but some of it simply takes the form of hiring "too little" or "too much" labor for the current level of production. These changes in employers' utilization of their existing workforce will show up as changes in labor productivity.
More capital-intensive production	This is a familiar story: When labor gets more expensive (or scarcer), businesses use more capital instead. This is usually what people mean when they say "Econ 101" shows that rising wages lead to less employment. This can be a problem when it's a question of rising wages due to legislation or unions, but it shouldn't be a problem when it's a question of rising wages due to labor scarcity. Insofar as businesses can substitute machines for labor, rising wages will not be passed on to prices, so there is more space to push unemployment down.
Innovation responds to demand and costs.	High wages not only lead businesses to adopt new technology, but also they can redirect resources toward innovation and developing new laborsaving techniques. We might first see these affects appear first at the firms facing higher labor costs and then diffused through the economy., Once developed, new technologies can be used by anyone.

¹⁵ One of the few attempts to measure this kind of disguised unemployment in modern economies is Eatwell (1997).

POSSIBLE LINKS FROM DEMAND TO PRODUCTIVITY

Employment shifts toward higher-productivity sectors	Aggregate productivity may rise because workers in a given industry produce more stuff. But it may also rise even if output per worker is constant in each individual industry, if there's a shift in employment toward higher-productivity industries. This may happen for one of two reasons: First, higher wages will disproportionately raise costs for more labor-intensive sectors. This will presumably depress growth to the benefit of less labor-intensive, more productive ones. Second, the more income-elastic sectors may also be higher-productivity ones. For example, when income falls, people are less likely to spend on consumer electronics, where productivity is high, than on personal services like haircuts, where productivity is relatively low.
Employment shifts toward higher-productivity firms	This sounds similar, but it's a different story. As Josef Schumpeter famously argued, technical change doesn't happen all at once across a whole society, but is normally embodied in new businesses founded by entrepreneurs. These new businesses normally need to bid away workers and other inputs from incumbent firms. Given the "frictions" in the labor market, this will require offering a wage significantly above the going rate. Meanwhile, the fact that the least productive firms can't afford to pay higher wages will cause them to decline or exit, which also raises average productivity. When wages are flat, on the other hand, low-productivity firms can continue operating. In this sense, higher wages are an integral part of productivity growth.
Strong demand favors innovation	This is similar to the previous story, but looks at the capital and product markets rather than the markets for labor and inputs. Productivity gains, again, disproportionately come from the emergence of new businesses, rather than from improvements at existing ones. It is much easier to start a new business in a boom than in a recession. Consumers are more receptive to new brands and new products when their incomes are high, and lenders and investors are much more willing to take risks in a boom than in a slump.
Increasing returns	This last story is in some ways the simplest. It may be that output per worker rises at the firm, industry or economy-wide level when the amount being produced increases. (This relationship is sometimes referred to as "Verdoorn's law.") It's worth noting that increasing returns is an area in which the intuition of people with economics training diverges sharply from people in the business and policy worlds. In many contexts outside economic theory, it is obvious that costs normally fall as more of something is produced. ¹⁶

All of these stories imply that higher demand should lead to higher measured labor productivity. Some of them—especially the first few—suggest a more or less immediate response of productivity to changes in demand, but also one that cannot continue indefinitely. These are best seen as forms of hysteresis, or reasons why short-term potential

¹⁶ For example, a recent article on Tesla notes that "a few years ago, Tesla broke ground on ... a facility in Nevada that will supply it with batteries on a massive scale ... economies of scale will make Tesla's batteries cheaper than those from third parties." To the author of this article and, presumably, to Tesla executives, it is obvious that as the number of batteries produced rises, the cost per battery will fall (Lee 2017).

may diverge from medium-term potential. They imply that a transient period—perhaps several years—of high demand and perhaps higher inflation could shift productivity and demographically adjusted employment back to their pre-recession trends, after which the space for demand policy to boost output would be exhausted.

Other mechanisms, especially the last two, are presumably slower. These will come into play only in the context of a sustained “high-pressure” (or “low-pressure”) economy, when demand pushes up against supply constraints (or falls short of them) consistently over a number of years. If these stories are important, then demand-boosting policy may indeed lead to rising wages, rising inflation, bottlenecks, and so on. (In the terms introduced earlier, the economy will be operating above short-run potential.) But if demand keeps pushing against supply constraints, eventually supply will adjust. If these later stories, focusing on innovation and pure returns to scale, are the important ones, then it is possible that strong demand can lead to faster output growth indefinitely. (The middle stories, focusing on investment and labor reallocation, presumably operate on intermediate time scales.)

Our economic problem is not that machines are replacing labor, thanks perhaps to wages that are too high. Rather, it is that machines are not replacing labor, because wages are too low.

If we follow the logic of the later stories a step further, we could even say that in the long run, our economic problem is not that machines are replacing labor, thanks perhaps to wages that are too high. Rather, it is that machines are *not* replacing labor, because wages are too low. An important strand of development economics has argued that the low wages of poor countries that have led to low productivity there, and not vice versa. Similarly, there’s a well-known argument that the reason the industrial revolution happened first in Britain rather than in China or India (or Italy or France) is not that that the necessary technical innovations were present only in Britain. They were present many places, but it was the uniquely high cost of British labor that made them profitable to adopt for production. But we needn’t accept these broad historical claims to think that the efficiency with which labor is used will depend, in part, on its price and availability. There are many reasons to think that a high-pressure economy, in which unemployment is low and wages are rising, will be more conducive to the various processes that raise output per worker.

Most importantly, right now we don’t need to know whether the link between demand and productivity is a short-run hysteresis phenomenon or based on the longer-run factors of innovation and pure returns to scale. Either way, there is considerable scope for continued

expansionary policy. The distinction between short-run hysteresis and a longer-run link between demand and productivity will only become important after a few years of a high-pressure economy, once the pre-2008 trend has been regained. At that point, hopefully, we will have more evidence as to whether higher demand can continue to boost productivity growth.

Longer-run demand effects on productivity are challenging for many economic models, in which spending decisions are made with respect to a start from a real endowment of resources and technology. But the shorter-run hysteresis stories, including effects of demand on productivity, are possible to incorporate into the DSGE models widely used in macroeconomics. Adding hysteresis to these models naturally leads to a prescription of more expansionary monetary and fiscal policy in response to downturns (Engler and Tervala 2016; Kienzler and Schmid 2014). For economists who are reluctant to accept claims that are not backed up with such formal models, this work may be reassuring.

3.2 THE SLOWDOWN PRODUCTIVITY GROWTH HAS COINCIDED WITH EXCEPTIONALLY WEAK INVESTMENT SPENDING.

Over the short run, labor productivity is likely to be procyclical because existing labor and capital are more fully used during booms, and underused during slumps. The first three stories in Section 3.1 are various ways of looking at this kind of link between productivity and demand. As we look at longer periods, this kind of straightforward underutilization—though by no means impossible—seems less likely. Over longer periods, the most straightforward link between demand and productivity is via investment spending. More investment means that each worker has more capital to work with. Since innovation is normally embodied in particular capital goods and has spillovers to other firms and industries, more investment also likely leads to higher output per worker even where measured capital-intensity does not rise. Conversely, weak investment implies more labor-intensive production techniques, and less productivity-boosting innovation. Economists sometimes try to distinguish between these two channels through the use of a production function, in which total factor productivity (TFP) represents the output of a given combination of workers and capital. It is not clear how meaningful such estimates are, since they depend on a number of unobservable parameters. Labor productivity is more straightforwardly measurable. But one can be agnostic about the value of the production-function technique and still see a clear link between investment and productivity growth.

In this light, a central fact about the current business cycle is the exceptionally weak growth in investment spending. Figure 13 shows peak-to-peak growth in real investment

spending in each of the past ten business cycles.¹⁷ As the figure shows, there is no precedent for the weakness of investment in the current cycle. Nearly ten years later, real investment spending remains less than 10 percent above its 2007 peak. This is slow even relative to the anemic pace of GDP growth, and extremely low by historical standards. In the three previous cycles lasting that long, real investment spending had increased anywhere from 30 to 80 percent. Even shorter cycles saw substantially greater investment growth.

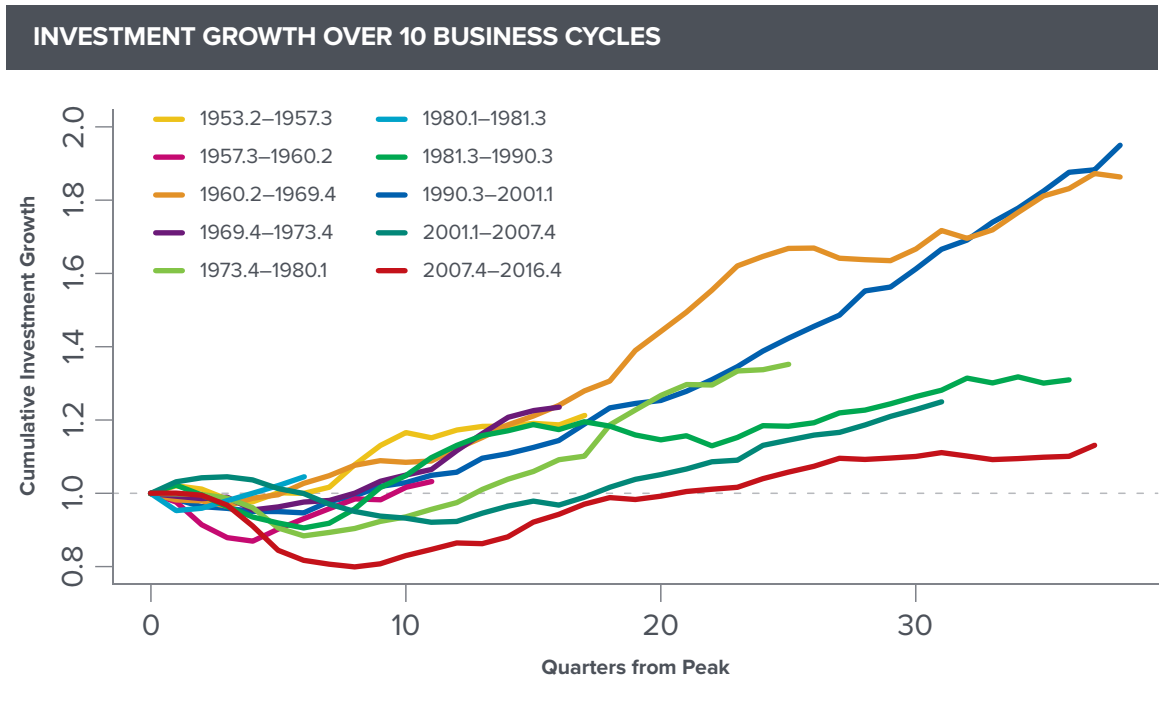


FIGURE 13 Cumulative real growth of private nonresidential fixed investment from business-cycle peak. *Source: BEA, author's analysis.*

As the figure shows, not only is investment growth weak in this cycle so far, it is getting even weaker. By the BEA's preferred measure, real investment spending peaked in the third quarter of 2015 and has declined slightly since then. And even this may overstate the real strength of investment. In 2013, the BEA made a major revision to the way it calculates investment spending, adding R&D and the production of intellectual property (movies, music, etc.) to the spending on plant, equipment and software that had previously been counted as investment. During the most recent cycle, business spending on IP production has been considerably stronger than other forms of investment. While there are good

¹⁷ The series is Private Nonresidential Fixed Investment, deflated by the BEA's implicit price deflator for private nonresidential fixed investment. Each series is set to 100 for the quarter designated a business cycle peak by the (National Bureau of Economic Research (NBER).

arguments for (and against) including this type of spending with investment, it is unlikely to have the productivity-boosting effects of more traditional forms of investment. If we limit our measure of investment spending to plant and equipment plus R&D, the current expansion looks even weaker at only 7 percent above the 2007 peak.

The link between investment spending and productivity growth is straightforward (if insufficiently emphasized in current debates). Could weak investment spending in turn be due, at least in part, to weak demand? Theories that link businesses' investment spending to their current sales are known as "accelerator" models of investment. The first modern accounts of business cycles emphasized accelerator dynamics, and they are still found in many forecasting models. But today's economic theory tends to downplay the accelerator mechanism. After all, the thinking goes, capital goods are long-lived, while demand-driven fluctuations are relatively brief. So business investment decisions should be more sensitive to factors like the expected return and the cost of capital, rather than on current demand.

But there has recently been a revival of interest in accelerator models. As Jason Furman notes, economists at the International Monetary Fund (IMF), OECD and the U.S. Council of Economic Advisors have found that "investment generally follows overall demand"—the path of current spending in an economy can explain most of the behavior of investment. This suggests there may be some merit in the older Keynesian view, in which investment mainly depends on whether current demand can be met with existing capacity; lower interest rates, more favorable tax treatment and similar factors are unlikely to have much effect as long as businesses can comfortably meet current demand with existing capital. As Furman observes, the simple demand story explains much, though not all, of the variation we see in investment growth: stories that "blame weak investment on pessimism and uncertainty among businesses, government regulations, limited access to credit, or ... taxation ... are harder to square with the data" (Furman 2015).

Weak investment growth offers a clear link between demand and productivity growth. (And if one accepts the accelerator story in general, slow investment growth then becomes an independent piece of evidence for continued weak demand.) If the accelerator mechanism is real, then central bankers and other policymakers who treat weak productivity growth as a supply constraint may be making a self-fulfilling prophecy. In the view of those who see the U.S. economy as currently close to potential, the slowdown in productivity growth represents a negative technology shock—an exogenous decline in the economy's productive capacity. In this view, the appropriate response of the Fed is to raise interest rates, to bring spending down to the new, lower path of potential output. But what if Furman is right, and weak investment is itself the result of weak demand, while low productivity growth in turn reflects weak investment? Then the Fed's contractionary stance will in fact be contributing to the productivity slowdown to which it believes it's responding.

3.3 THE PRODUCTIVITY GROWTH SLOWDOWN IS WIDELY SHARED ACROSS SECTORS, EVEN LESS TECHNOLOGY-INTENSIVE ONES.

One common argument for seeing the productivity slowdown as a negative technology shock is that it is simply the flipped of the positive shock of the “tech boom” of the 1990s. In this view, the fact that the slowdown in productivity was not predicted before the recession should not be seen as evidence that it is the result of the recession. Rather, it is simply the correction of an earlier, mistaken extrapolation from a period of exceptional productivity growth. Probably the most influential statement of this view is found in Gordon (2016), but it is widely held among policymakers. For example, John Fernald of the San Francisco Fed describes the productivity slowdown as follows: “Information technology fueled a surge in U.S. productivity growth in the late 1990s and early 2000s. However, this rapid pace proved to be temporary, as productivity growth slowed before the Great Recession” (Fernald 2015).¹⁸ In this view, slow investment growth presumably reflects lower expected profits, rather than weak demand.

One problem with this view is that it requires an exceptionally unlucky coincidence. The end of the housing bubble, the financial crisis, and the related fall in spending by households and businesses are a demand shock—a change in desired spending. Technological exhaustion, as postulated by writers like Gordon and Fernald, is a supply shock—a change in the economy’s productive potential. To frame the post-2008 slowdown in productivity growth as the result of autonomous technological factors is to say that the worst demand shock in more than 60 years was immediately followed by the worst supply shock in more than 50 years, by sheer chance. (See Blanchard, Cerutti, and Summers 2015 on this point.)

Still, unfortunate coincidences do occur. A stronger objection to the technological-exhaustion view is that it seems inconsistent with the larger macroeconomic picture. A negative supply shock is normally associated with higher inflation as well as lower output growth. And if the return on capital has fallen, that should mean a lower share of income flowing to capital owners. But as discussed in Section 4.1 below, in fact we see the opposite.

A second problem with the information-technology story is that labor productivity growth is low not just relative to the late 1990s and early 2000s, but also to any other

¹⁸ Fernald continues: “Furthermore, looking through the effects of the economic downturn on productivity, the reduced pace of productivity gains has continued and suggests that average future output growth will likely be relatively slow.” This is a good example of the danger the present paper is warning against, that the mere persistence of slow growth will be taken as evidence that it is a supply-side phenomenon out of reach of macroeconomic policy. In this way, a failure to solve demand problems in the past can become an argument against trying to solve them in the future.

extended period in postwar history. And, critically, it is not particularly concentrated in the information-technology sectors, but is widely distributed across industries—including many that did not share in the late-1990s boom.

In the U.S., as in other advanced countries, productivity dynamics are normally dominated by a handful of sectors. Most productivity gains are accounted for by a relatively small number of capital-intensive industries, while others—many services, but also construction—show little or no productivity improvement over long periods. The tech boom of the late 1990s and early 2000s, with its rapid productivity gains in the broadly defined information sectors, was a typical example of this. If the productivity slowdown of the past decade represented an exhaustion of technological progress, one would expect the slowdown to be concentrated in the same industries that have most benefited from technological progress historically. This is the implicit hypothesis behind claims that the overestimate of future productivity gains prior to the recession simply represented a misguided extrapolation from the tech boom period. But when we turn to the data, this is not really what we see.

It is immediately clear that, while the high-tech sectors do account for a disproportionate share of the productivity slowdown, it is not limited to these sectors. On the contrary, what's most striking about the productivity slowdown is how widespread it is.

The BEA publishes productivity data for 67 disaggregated industries. To avoid the dynamics of the Great Recession itself, we'll look at the period 2011-2015 (2016 industry-level data is not yet available), compared with the 2001-2007 expansion. It is immediately clear that, while the high-tech sectors do account for a disproportionate share of the productivity slowdown, it is not limited to these sectors. On the contrary, what's most striking about the productivity slowdown is how widespread it is. Compared with 2001-2007, productivity growth in 2011-2015 was 1.6 points per year slower, essentially all of which is slower within-industry growth. The high-tech sector accounts for about 0.65 points per year of this, or about 40 percent.¹⁹ But the slowdown in productivity is remarkably widespread. Over 80 percent of the industries broken out by the BEA show slower productivity growth in the second period than in the first.

¹⁹ Here and in Figure 14, we define the high technology industries as broadcasting and telecommunications, computer and electronic products, computer systems design and related services, and information and data processing services.

PRODUCTIVITY GROWTH BY PERIOD IN TWO PERIODS

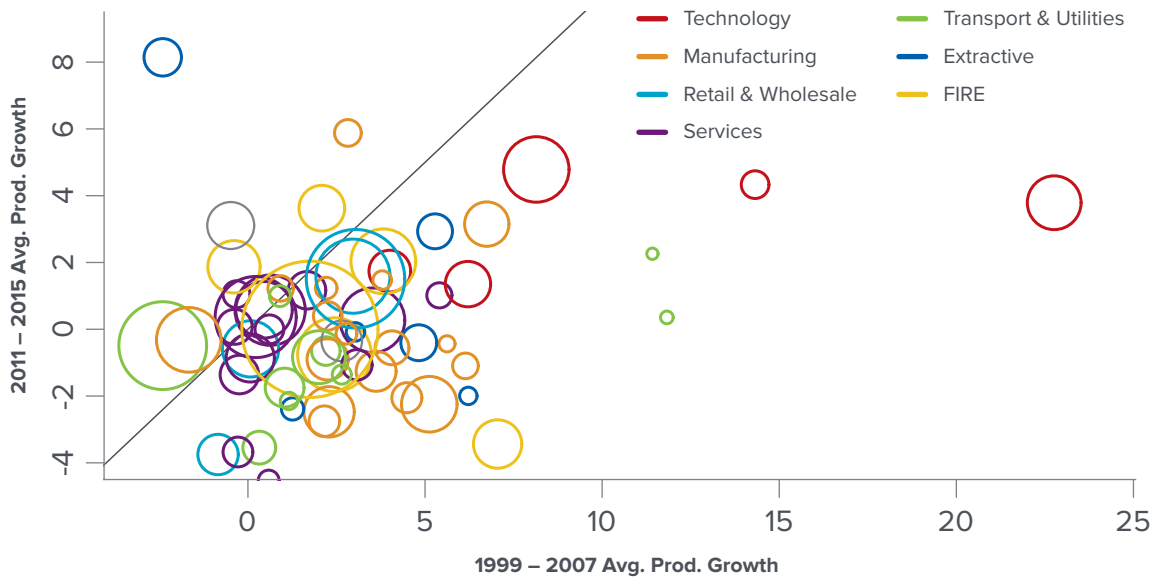


FIGURE 14 Productivity growth by industry. The horizontal axis shows average annual real productivity growth over 1999-2007, the vertical axis shows average annual real productivity growth over 2011-2015. The size of the circles is proportional to the average employment in the industry over the full period. *Source: BEA, author's analysis.*

This is shown in Figure 14. For each industry, productivity growth in the first period is shown on the horizontal axis; productivity growth in the second period on the vertical. The diagonal line shows equal rates of productivity growth in the two periods. The points are colored by sector, and scaled by employment. As is clear, the vast majority of industries, representing the vast majority of employment, showed slower productivity growth in the second period than in the first. If the main story was the end of the 1990s tech boom, and over-optimistic extrapolation from that period, then the largest declines in productivity should be found in the industries that made the biggest gains in the earlier period.

Graphically, the slope of the line would be flatter than 45 degrees. But that is not what we see. Rather, the main change is a downward shift in productivity growth. For the most part, productivity in industries that showed modest productivity gains in the earlier period decelerates just as much as in those that saw faster gains.

It is true that the largest slowdowns took place in the tech sectors—the three large red circles to the right of Figure 14 are, from left to right, broadcasting, information and computer manufacturing. But the productivity slowdown is far from limited to these industries—the great majority of points in the figure fall below the diagonal line. Again,

about 40 percent of the overall decline in productivity growth is accounted for by sectors where the declining return on information technology is plausibly a big part of the story. But most of the decline is accounted for by sectors like real estate, retail and wholesale trade, food and beverages, and transportation. It is hard to explain the productivity slowdown in these sectors in that way.

3.4 PRODUCTIVITY AND THE BUSINESS CYCLE

Besides the widespread nature of the productivity slowdown, there is a second, more subtle, but important, line of evidence suggesting that the slowdown may be due to weak demand. This is precisely the fact that the past decade has seen large declines in both productivity and employment growth. Both theory and history suggest that extended periods of below-average productivity growth are normally periods of above-average employment growth, and vice versa. During business cycles, on the other hand, employment and productivity move together. So the fact that both productivity and employment have grown slowly since 2007 makes the last decade resemble an extended recession, more than the kind of medium-run shocks the economy has experienced in the past.

The fact that both productivity and employment have grown slowly since 2007 makes the last decade resemble an extended recession.

In the supply side story, the slow labor force growth of the past decade is due to demographic and other social factors reducing the supply of available labor, while the productivity slowdown is independently due to a decline in technological progress. Both are supposed to be exogenous changes, independent of the state of demand. But a little thought suggests that this story is harder to sustain than it seems at first glance. (We will return to this point in Section 4.1.) Over the medium term, we would expect these two components of output growth to move inversely, not together.

Both common sense and economic theory suggest that sustained periods of slow labor force growth should see an acceleration of labor productivity. A reduced supply of workers means that labor is scarcer; this implies rising wages and a greater return on labor-saving innovations. It also means that any scarce resources are available in greater amounts per worker. As a result, we would expect to see an exogenous slowdown in labor force growth associated with more rapid productivity growth.

Conversely, technological change is normally laborsaving. (This doesn't have to be true as a matter of logic, but it is a very clear general pattern from the past 200-plus years of economic growth.) Slower technological progress should mean more labor-intensive production, which means more employment for a given level of demand. In other contexts, we often think of technology as reducing jobs.²⁰ Logically, if faster technological progress reduces employment, slower progress should increase it. So an exogenous slowdown in productivity growth should be associated with more rapid employment growth.

And indeed, for the U.S. and many other advanced countries, this is exactly what we see. Over periods of more than a few years, faster productivity growth is associated with slower employment growth, and vice versa.

But over shorter periods, the opposite pattern appears. Around business cycle turning points in particular—recessions and recoveries—productivity and employment growth are normally in sync. The fact that output, employment and productivity move together over the business cycle has been recognized by empirically oriented macroeconomists since the early 20th century. The fact that productivity falls during recessions means that employment varies less over the cycle than output does.

This behavior is quite stable over time, giving rise to Okun's law, an empirical regularity that is a standard tool of macroeconomic forecasting. In the U.S., Okun's law says that the unemployment rate will rise by one point for each 2.5-point shortfall of GDP growth over trend—a ratio that doesn't seem to have changed much since Arthur Okun first described it in the mid-1960s (Okun 1962).

The different relationships between employment and productivity growth over different time periods are shown in Figure 15.

Okun's Law: The unemployment rate will rise by one point for each 2.5-point shortfall of GDP growth over trend.

²⁰ See the many predictions of automation, robots replacing human labor, etc.

REGRESSION OF EMPLOYMENT GROWTH ON PRODUCTIVITY GROWTH




FIGURE 15 Regression coefficient of employment growth on productivity growth. This shows the coefficient of a regression of the employment growth rate on the productivity growth rate over periods ranging from one quarter up to 10 years. The horizontal axis shows the number of quarters over which the change is measured, the vertical axis shows the coefficient. Source: BEA, author's analysis.

The figure shows the coefficient of a regression of employment growth on productivity growth over various periods, from one quarter out to 40 quarters. (This is based on postwar U.S. data.) Over periods of a year or less, the correlation is positive. Six-month or nine-month periods in which employment growth was unusually weak are likely to have seen weak productivity growth as well. This is the cyclical effect discussed above. But over longer periods, autonomous variation in productivity and labor supply are more important, and these tend to produce a negative relationship between employment and productivity. And that's exactly what we see. A ten-year period in which productivity grew unusually quickly is very likely to be one in which employment grew slowly, and vice versa.

The period since 2007 has seen close to the slowest employment growth of the postwar period and close to the slowest productivity growth.

This relationship was quite stable in the 60 years between 1947 and 2007. But in the past decade it has completely broken down. The period since 2007 has seen close to the slowest employment growth of the postwar period and close to the slowest productivity growth.

It is normal for employment and productivity to move together for a couple quarters or a year, but very unusual for this joint movement to be sustained over nearly a decade. In the postwar U.S., at least, periods of slow employment growth are much more often periods of rapid productivity growth, and vice versa.

To put it another way: The past decade has seen exceptionally slow growth in employment—about 5 percent over the full period. If you looked at the U.S. postwar data, you would predict with a fair degree of confidence that a period of such slow employment growth would see above-average productivity growth. But in fact, the past decade has also seen very low productivity growth. The relation between the two variables has been much closer to what we would predict by extrapolating their relationships over periods of a year. In that sense, the current slowdown resembles an extended recession more than it does previous periods of slower growth. And since no one disputes that the demand side drives recessions and recoveries, this creates a prima facie case that the weakness of the current expansion may be demand driven as well.

The current slowdown resembles an extended recession more than it does previous periods of slower growth.

As we'll see in the next section, the macroeconomic picture looks even more like an extended period of weak demand when we look at additional variables such as inflation, wages and interest rates.

The big picture: Overheating?

4.1 THE OVERALL STATE OF THE ECONOMY IS HARD TO RECONCILE WITH A STORY IN WHICH WE ARE FACING SUPPLY CONSTRAINTS.

In the conventional story, productivity is supposed to be driven by technology, so a slowdown in productivity growth reflects a decline in innovation and so on (Gordon 2012). Employment is driven by demographics, so slower employment growth reflects aging and small families. Both of these developments are negative shifts in aggregate supply, so they should be inflationary. If the economy's productive potential declines, then the same growth in demand will instead lead to higher prices. To maintain stable prices in the face of these two negative supply shocks, a central bank would have to raise interest rates in order to reduce aggregate spending to the new, lower level of potential output. This is not what we have seen. We have seen declining inflation even as interest rates are at historically low levels. So even if you explain slower productivity growth by technology and explain slower employment growth by demographics, you still need to postulate some large additional negative shift in demand.

What are the major macroeconomic developments of the past ten years compared with the previous fifty? For the U.S. and most other developed countries, the list looks like:

- Low and falling inflation
- Low and falling interest rates
- Slower growth of output
- Slower growth of employment
- Low business investment
- Slower growth of labor productivity growth
- A declining share of wages in income

Any macroeconomic story has to make sense of the whole picture. And while it's easy to think of stories about the economy's underlying productive capacity that work for individual items, it's practically impossible to explain the whole combination except through weak demand.

Let's think about this in the aggregate supply-aggregate demand framework used by economics textbooks and most macroeconomic policymakers. In that framework, low inflation must be the result of a negative demand shock, a positive supply shock or a shift toward more contractionary policy—higher interest rates or the equivalent. Setting aside the demand story for a moment, what do we have?

The textbook says that, over the long run, interest rates must reflect the marginal product of capital—the central bank (and monetary factors in general) can only change interest rates in the short run, not over a decade or more. Consequently, a sustained decline in interest rates must mean a decline in the marginal product of capital. This is consistent with slower productivity growth. But there is a problem. The return on capital should also reflect its marginal product, so a declining marginal product of capital should imply a lower share of profits in national income, and a rising share of wages—just the opposite of what we've seen.²¹

Another problem is that a declining marginal product of capital is a negative technological shock. And negative technological shocks should be associated with rising inflation—if the productive capacity of the economy grows more slowly, the same flow of spending will bid up the prices of existing goods rather than call forth new production. But rising inflation is, again, the opposite of what we've actually seen.

Employment growth in this framework, meanwhile, is determined by demographics, or perhaps by structural changes in labor markets that change the effective labor supply. Slower employment growth means a falling labor supply—and that should, again, be inflationary, and it should be associated with higher wages. If labor is becoming relatively scarce, its price should rise. It doesn't matter whether we think of wages as being set by the marginal product of labor or by a bargaining process between workers and employers, since in this case, the two stories agree. A fall in the labor supply will increase the bargaining power of workers, and it will also raise the marginal product of labor as it's withdrawn from the least productive activities—that's what "marginal" means. Either way, the demographic story of falling employment is inconsistent with low inflation, with a falling wage share and the slowdown in productivity growth. But until recently, the slowdown in employment growth was not accompanied by a rise in wages or the wage share. And broader measures of labor market conditions suggested that labor markets were still weak (Pacitti and Fichera, 2015; Pacitti, 2017).

It is true that in the past year, the wage share has risen. But this increase is modest, especially given the headline unemployment rate, and not nearly enough to offset the

²¹ An increase in market power could help explain the combination of low interest rates, high profit rates and weak investment. Profitable companies with access to cheap finance may nonetheless decline to expand if they enjoy stable monopolies in their markets. There is a strong case for seeing demand-boosting policy and strengthened anti-trust enforcement as complements (Konczal and Steinbaum 2016).

declines in previous years. Indeed, while real hourly compensation did increase by 2.7 percent in 2015—the fastest rate in nearly 20 years—this followed an extended period (2008 to 2013) in which there was no increase in real hourly compensation at all. And in 2016, real compensation growth fell back to a more modest 1.2 percent. So the broader wage picture still looks like one of weak labor demand. Compensation may accelerate again, and other labor-market measures—like the number of job openings and the rate of voluntary quits—do suggest a tighter labor market. But even if this was the case in the past year or so, we still have to explain the slow growth in the previous five years.

Both weak labor markets and slower growth of labor productivity might be explained by an increase in labor supply, but then why has employment decelerated so sharply? More often, productivity is taken as technologically determined. Slower productivity growth then implies a slowdown in innovation—which at least is consistent with low interest rates and low investment. But this “negative technology shock” should, again, be inflationary. And it should be associated with a fall in the return to capital, not a rise.

On the other hand, the decline in the labor share is often understood to reflect a change in productive technology that encourages substitution of capital for labor. For example, the IMF’s most recent World Economic Outlook attributes about half of the long-term fall in the labor share to technological progress, particularly in information and telecommunications, that allowed many jobs to be automated (IMF 2017). We don’t take a position here on whether this is a reasonable way to think about the longer-term trend in factor shares. But it faces obvious problems as applied to the past decade in the U.S. (and other countries with similar experiences). After all, this is the same period that writers like Gordon (2016) and Fernald (2015) say has seen a decline in productivity growth thanks to the exhaustion of information technology. In isolation, stories that explain a falling wage share by rapid technological progress, and falling productivity growth by slow technological progress, may each seem plausible. But both cannot be true simultaneously. We also may wonder how an explanation of falling wage shares in terms of capital deepening is reconciled with the fall in interest rates, and with weak investment. To replace workers with robots, someone has to make the robots, and someone has to buy them. And again, by definition, this raises the productivity of the remaining workers.

Excessively contractionary policy by the central bank could explain much of the phenomena above. But there’s an obvious problem. The interest rate set by the Fed has been at zero for the past decade, and the Fed has taken unconventional expansionary policy as well. Yes, you could certainly argue—as many people have—that even this exceptionally expansionary stance was not expansionary enough. But then you still need to postulate an exceptionally deep and persistent negative demand shock, to explain why such extremely loose monetary policy was needed.

As stated above, it is not hard to tell a supply-side story that does a good job explaining one or a couple of the phenomena listed above. Much of current macroeconomic policy debate relies on a mix of disconnected stories that fit a couple pieces but are contracted by other pieces, and by each other. (On Mondays and Wednesdays, wages are low because technological progress has slowed down, holding down labor productivity. On Tuesdays and Thursdays, wages are low because technological progress has sped up, substituting capital for labor.) But it is very hard to tell a supply-side story that is consistent with the whole set at once.

Another problem with the supply-side story is the fact that the slowdown happened immediately after the deepest recession in postwar U.S. history. Almost no one disputes that the recession itself was driven by demand.²² So it would be quite a coincidence if the deepest fall in aggregate demand since the Great Depression just happened to be followed by an unrelated collapse in aggregate supply.

To believe the supply-side story, you need not just one but a series of coincidences. You have to believe that just as the economy recovered from the financial crisis, there was a steep slowdown in employment driven by demographic change. And at the same time, there was an unrelated slowdown in the pace of technological innovation. Even then, you still need an independent fall in demand to explain why the negative shocks to labor and technology did not produce the inflation that the textbook predicts. And the timing has to be perfect: if the demographic decline in the labor force had come a few years earlier, there would have been no rise in unemployment during the recession, while if it had had come a few years later, we would still be looking at high unemployment today. Since we need a fall in demand in any case, Occam's razor suggests we should prefer stories in which that drives everything else: a weak labor market discourages labor force participation, while weak demand in product markets and stagnant wages together discourage productivity-boosting innovation. Rather than a precisely synchronized series of unobservable shocks, this story requires just one, which we have in fact observed in the form of the financial crisis.

The need to look at the full macroeconomic picture to distinguish supply from demand was expressed forcefully by DeLong and Summers (2012). They argue that interpretations of the fall in GDP as a decline in potential are missing an “elementary signal extraction point”: a fall in potential output implies a simultaneous fall in production and increase in inflation. Just as a simultaneous fall in the price and quantity sold of a given good can't be explained by a fall in supply of it but only by a fall in demand for it, a simultaneous slowdown in production and prices in the economy as a whole suggests a fall in aggregate demand, not in aggregate supply.

²² Exceptions include a few economists at the University of Chicago and similar “freshwater” institutions, who blame the downturn on a fall in productive potential due to increased regulation and government spending. For a good example, see Casey Mulligan, *The Redistribution Recession*. But this is very much a minority position.

Consequently, we shouldn't base our judgments about potential output merely on GDP itself, or on its components (employment and productivity). We need to look at other macroeconomic variables as well, such as inflation. A decline in the economy's productive potential will be associated not just with slower output growth, but with higher inflation and with rising interest rates—the textbook response to a negative supply shock. When we see a simultaneous fall in output, inflation and interest rates, we are looking at a situation of inadequate demand.

This point is illustrated in Figure 16, which shows the percentage change in productivity, employment and prices over every ten-year period from 1947-1957 through 2006-2016. The colors show inflation—periods with moderate inflation are gray, periods with unusually high inflation are yellow and periods with unusually low inflation are blue. Growth between 1990 and 2000 is, for instance, represented by the point labeled “2000”. During this decade, total employment rose by about 20 percent while productivity rose by a total of 15 percent, implying an annual real growth of 3.3 percent, very close to the long-run average.

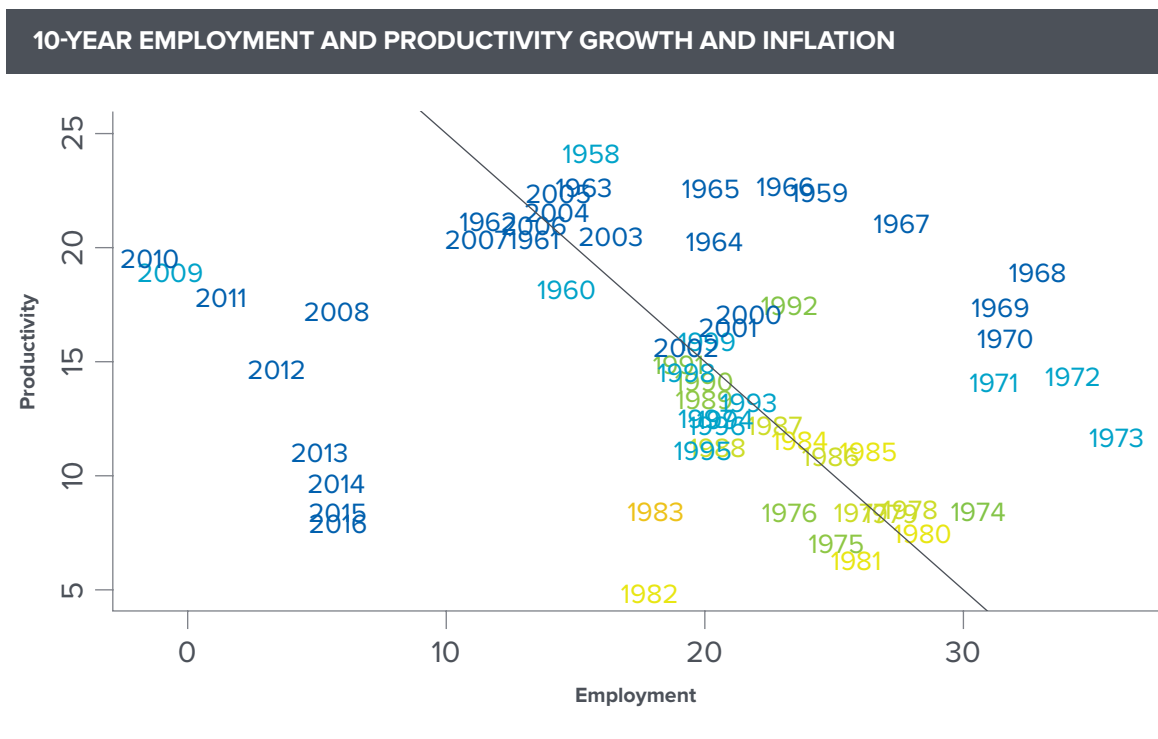


FIGURE 16 Employment growth, productivity growth and inflation over rolling 10-year periods. The horizontal axis shows the average rate of employment growth and the vertical axis shows the average rate of labor productivity growth, both for the decade ending in the year shown. The color of the text corresponds to the average inflation rate over the period—yellow is high inflation, blue is low. The diagonal is a regression line *Source: BEA, author's analysis.*

As we can see, most of the periods are distributed along the main diagonal running from upper left to lower right. It's easy to interpret these points in terms of supply side stories. Toward the upper left are periods of slow labor force growth and rapid productivity growth; these will be times when slow labor force growth was encouraging laborsaving innovations, and/or when rapid innovation meant that automation displaced jobs more rapidly than they could be replaced. In the lower right are periods of rapid employment growth and slow productivity growth; these are times when rapid growth of the labor force and/or slower innovation shifted the mix of jobs toward more labor-intensive, less productive activities. As the figure shows, most of the medium-term macroeconomic variation in the postwar U.S. can be understood in these terms, but there are two exceptions.

The first is the 1970s, when both employment and output per worker rose rapidly. To some extent, this may have reflected a combination of large numbers of young people entering the labor force and the dissemination of various new technologies. But as the colors show, this was also a period of high inflation, and we also know that it was the period during which the labor share of output reached its peak, and when the Fed, eventually, felt compelled to sharply raise rates. As a result, it is natural to see this as a period of exceptionally strong demand: a “high-pressure economy,” in the language of the time. The other exception is, of course, the current period.

In other words, yellow points above the line suggest positive demand shocks—an increase in spending can increase output growth above trend, at least for a while, but will pull up prices as well. Alternatively, we can think of yellow points below and to the right of the line as suggesting negative supply shocks. If the productive capacity of the economy declines for some reason, output growth will slow, and prices will rise as private actors—abetted by a slow-to-react central bank—attempt to increase spending at the usual rate. Similarly, blue points above the line suggest positive supply shocks. And blue points below the line suggest negative demand shocks. As DeLong and Summers say, if a fall in employment and output is accompanied by low or falling inflation, it probably reflects a fall in desired spending, not in the economy's productive potential. Yellow points below the line suggest negative supply shocks—the best case for this is the decade 1973-1983.

We immediately see what an outlier the recent period is. Both employment and productivity growth over the past ten years have been dramatically slower than over the preceding decade, at about 5 percent each, down from about 20 percent. The ten-year periods from 2000-2010 and 2001-2011 were the only such examples in postwar U.S. history during which total employment actually declined. The abruptness of the deceleration on both dimensions is a challenge to views that slower growth is the result of deep structural forces. And the combination of the slowdown in output growth with falling prices—especially given ultra-low interest rates—strongly suggests that we've seen

a negative shift in desired spending (demand) rather than in the economy's productive capacities (supply).

Another way of looking at this is as three different regimes. In the middle is what we might call “the main sequence.” Here there is steady growth in demand, met by varying mixes of employment and productivity growth. On the upper right is what gets called a “high-pressure economy,” in which low unemployment and strong demand draw more people into employment and facilitate the reallocation of labor and other resources toward more productive activity, but put upward pressure on prices. On the lower left is stagnation, where weak demand discourages participation in the labor force and reduces productivity growth by holding back investment, new business formation and by leaving a larger number of those with jobs underemployed, and persistent slack leads to downward pressure on prices (though so far not outright deflation). From the perspective of macroeconomics, the past decade has been a sort of anti-1960s.

4.2 POTENTIAL IS NOT A HARD LINE; POLICY NEEDS TO TAKE INTO ACCOUNT NOT ONLY WHERE IT IS, BUT THE CONSEQUENCES OF GUESSING WRONG.

A final and important problem is that it is not enough to know where the full employment or potential output line is. Policymakers also need to consider the risks of overestimating potential output versus underestimating it. The costs of inadequate demand are clear: in the short run, unemployment and useful goods and services not produced. And in the long run, if you accept some version of hysteresis, there is a diminished labor force and less productive businesses. Under current conditions, though, the costs of excessive demand are less obvious. We normally think of these costs in terms of “overheating” or “crowding out,” but what do these terms concretely mean?

Policymakers also need to consider the risks of overestimating potential output versus underestimating it.

During the 1960s and 1970s, it was widely believed that central banks could trade off unemployment against inflation, accepting more of one in return for less of the other. The idea of that sort of tradeoff has become deeply unfashionable in recent decades. But in today's conditions, this kind of reasoning may again be relevant. As Blanchard (2016) argues, structural changes in the U.S. economy have weakened the relationship between

unemployment and inflation; it no longer seems to be the case that when unemployment falls below the NAIRU level, that will reliably lead to a substantial rise in inflation. So while historically it may have been the case that the goals of stabilizing inflation and stabilizing unemployment lined up, today “the Fed faces a tradeoff between stabilizing unemployment and stabilizing inflation,” and there is no reason to think that hitting its price stability target perfectly is the socially optimal policy. If lowering unemployment by one point would raise inflation by only 0.2 points, as Blanchard estimates, that presents “a very attractive tradeoff between inflation and unemployment.” Given the existence of hysteresis, there is an argument “to lower unemployment below the natural rate for some time.”

While historically it may have been the case that the goals of stabilizing inflation and stabilizing unemployment lined up, today “the Fed faces a tradeoff between stabilizing unemployment and stabilizing inflation,” and there is no reason to think that hitting its price stability target perfectly is the socially optimal policy.

The effects of more expansionary policy that are normally considered costs may in fact be benefits in today’s conditions. It is incumbent on FOMC members and other policymakers who support contractionary policy to explicitly state what they believe the costs of overheating are. Ten or twenty years ago, the fact that the economy was at full employment could be sufficient justification for a shift toward more restrictive policy. But today that is not enough; it’s also necessary to explain why overfull employment is undesirable.

So why do many policymakers think that once full employment is reached, further expansionary policy is useless or harmful? The simplest answer is that, as long as there are unused resources in the economy, increased spending will put them to work. But once there are no more unemployed workers, idle machines or empty buildings, additional dollars will instead bid up the price of the goods already being produced. In this story, overheating just means higher inflation.

A more sophisticated version of the story focuses on the labor market specifically. In this story, found in the standard “three equation models” used in advanced undergraduate textbooks and macroeconomic forecasting, nominal wage growth is determined in a formal or informal bargaining process in which the strength of the two parties depends on the unemployment rate. When unemployment is low, workers have many alternatives and

can demand faster wage increases. When unemployment is high, the advantage is on the employers' side, and wage growth will be slow. The goal for the central bank, then, is to keep unemployment at the level where nominal wage gains are just equal to productivity growth plus the target inflation rate. In this version, overheating means faster wage growth, which, presumably, will lead to higher inflation.

A third version of the story is that as demand rises past the limits of what the economy can produce, interest rates will rise. In older versions of this story, like the well-known IS-LM of Hicks, this interest rate rise happened automatically as credit reached the limits of the liquidity available to the economy. In more recent versions, interest rates rise because the central bank increases them once the economy reaches potential. Either way, excess spending by government (or consumption spending by households) "crowds out" productive investment.

So we have to ask: Are these stories still valid under current conditions? Would the consequences of too much demand (higher inflation, faster wage growth, higher interest rates) really be bad things under current conditions? And if so, how do we weigh their costs against the other side of inadequate demand?

Some form of the first story must be true, by definition. We can see this using a straightforward accounting identity:

$$(3) \text{ Percent change in nominal spending} = \text{percent change in real output} + \text{inflation rate} = \\ \text{percent growth in employment} + \text{percent change in productivity} + \text{inflation rate}$$

A dollar of additional spending on goods and services must result either in additional goods and services being produced, or a higher price of the ones already being produced. There is nowhere else for the money to go. And an increase in the quantity of goods and services being produced can only come through more workers being employed, more being produced by each worker, or some combination of the two. So if we think that the current unemployment rate means that no more workers are available, and that productivity increases are independent of demand, at this point further spending increases can only mean faster inflation.

As a matter of logic, this story is inarguable,²³ but it raises two questions. First does the current unemployment rate in fact mean that higher spending cannot call forth higher real output? The arguments in Sections 1 and 2 suggest not. There is good reason to think that with

²³ With one caveat: It is strictly true only for a closed economy. For a discussion of whether foreign trade creates a constraint for expansionary policy, see Mason 2016, "Dealing with the Trade Deficit".

stronger demand, labor force participation and labor productivity would rise. If we accept one or both of those claims, then there is additional output and—in the first case—possible employment that we won’t see without some intervention to raise the level of spending in the economy. But even if we are skeptical about those stories, we still have to ask how serious a cost higher inflation is at this moment. By most measures, inflation remains below the 2 percent target set by the Fed. If policy shifts rapidly in a contractionary direction, there is no guarantee that the target will ever be reached, especially if something else—a recession or negative shock—reduces private spending in the near future. It is hard to argue that the danger of higher inflation is an argument for contractionary policy when, as Larry Summers observed recently in *The Washington Post*, “our problem today is insufficient inflation.”

Many economists have suggested that a higher target—3 or 4 percent—would give the central bank more space to lower rates when, as is inevitable, more expansionary policy is needed in the future.

There is also less consensus that the 2 percent target is appropriate. Many economists have suggested that a higher target—3 or 4 percent—would give the central bank more space to lower rates when, as is inevitable, more expansionary policy is needed in the future. There is little dispute that the Fed’s inability to reduce short-term interest rates below zero, as its standard policy rule would have called for, limited its ability to counteract the fall in demand in the 2008-2009 recession. This meant it did less to boost demand than in previous recessions, leading to higher and more prolonged unemployment and a greater loss of output than would likely have occurred if it had been possible to lower rates further. While unconventional policy—quantitative easing, forward guidance, etc.—is intended to overcome this limitation, it is unclear how effective it is, and it carries costs of its own. A simpler solution is to allow higher inflation in normal times, so the Fed has more space to move real rates into negative territory when it needs to. Other economists support targeting the “level path” of prices rather than the inflation rate, or in other words trying to make up for inflation shortfalls or overshoots with deviations in the other direction; or support targeting the path of nominal GDP, or NGDP, rather than the price level. Both of these alternative policy rules would imply that the Fed should be targeting a higher inflation rate today.²⁴

²⁴ For example, see Binder and Rodrigue’s recent argument that the Fed’s current policy rule is insufficiently focused on maintaining full employment, and that an NGDP target would perform better. Note they see this as a medium-term goal that could supplement or replace the current inflation target, rather than a fixed rule that could be implemented automatically through futures markets, as proposed by some advocates of NGDP targeting.

There is also the question of balance of risks. While some people worry about runaway inflation if the Fed overshoots its target, the empirical evidence strongly supports the idea that expectations are well anchored—that a small rise in inflation would not lead people to expect further rises. And if it does so, it is straightforward to rein it in through a contractionary shift once it becomes a problem. On the other hand, it is entirely possible that the economy will face further deflationary pressures in the coming years. In this case, a policy that appears cautious today could end up being overtight in retrospect. It's a question of which scenario is more likely, which would be more costly and which would be harder to correct if the Fed does turn out to be wrong.

If we do want a higher inflation target, there is no way to reach it without a period of overfull employment.

Obviously, these are controversial positions often involving technical arguments. Our claim here is not that supporters of an alternative inflation target are right; it is that they are not so obviously wrong that the current target can be taken for granted. Supporters of the existing inflation target need to make an affirmative case for it, rather than just asserting that any inflation over 2 percent is unacceptable. We also want to make a slightly more subtle point: If we do want a higher inflation target, there is no way to reach it without a period of overfull employment. The relationship described by Equation 4 works both ways: just as spending beyond the economy's productive capacity raises inflation, you can't get higher inflation without spending beyond the economy's productive potential. So it makes no sense to suggest that a higher inflation target would be desirable, and simultaneously oppose further expansionary policy on the grounds that the economy is already at capacity.

In reality, of course, potential output is not a hard line, after which all prices rise in tandem. As the three-equation model highlights, prices of scarce factors will rise first; wages, in particular, are disproportionately sensitive to business cycle conditions. In the world of economics textbooks, these wage increases are passed on one for one to price increases. In reality, though, there is no reason to think that all wage increases will be passed on to prices; after all, firms have other costs, and there may be constraints on their ability to raise prices. This seems especially true in recent years, when the correlation between wage and price increases has gotten weaker, but it probably never was that close, so we need to consider other possible outcomes.

What happens when wages increase? We can answer this question in a general way with a second accounting identity:

$$(4) \text{ Percent change in nominal wages} = \text{percent change in productivity} + \text{percent change in labor share} + \text{inflation rate}$$

An increase in wages must result in some combination of higher output per worker, higher prices, or a greater share of output going to wages. Together, these three will always add up to the change in wages.

In textbooks, it is normally assumed that productivity is determined by technology, while wage and profit shares are fixed. So any increase in wage growth beyond the given rate of productivity growth must be passed on to inflation. But in the real world, things may not be so simple.

Higher wages may instead be accompanied by faster productivity growth, as in one of the stories in Section 3.3. Alternatively, wages may rise without a corresponding increase in either prices or productivity. Imagine that workers' wages rise but output doesn't change, and neither does the price of the finished goods. Since a business's revenue is unchanged but workers' income is higher, the income going to capital must be smaller. In other words, the wage share of income rises.

Equation (4) is true by definition. It always holds. It's important to understand this, to avoid confusion. People sometimes say that the central bank should ensure that real wages don't rise faster than productivity growth, or that nominal wages don't rise faster than productivity growth plus the target rate of inflation. Both of these formulas sound reasonable, but it's important to realize that they are both the same as saying that the central bank should ensure the wage share of income does not rise. This sounds much more problematic, but again, it is exactly equivalent to the two other expressions.

The idea that macroeconomic policy affects the distribution of income follows directly from standard macroeconomics. Let's say that we believe, as both common sense and the textbooks tell us, that the rate of wage growth depends on the level of unemployment. Suppose we define full employment in the conventional way, as the level of unemployment that leads to nominal wage growth just equal to productivity growth plus the central bank's inflation target. Then, by definition, any increase in the wage share requires a period of overfull employment, or unemployment below the full employment level. This holds even if you think the labor share in the long run is entirely technologically determined. There still must be transition periods in which wages are rising more rapidly than productivity growth, before the new long-run equilibrium is reached. And this in turn requires a period of overfull employment.

Over the past 15 years, the wage share in U.S. national income has fallen from 63 to 57 percent. (See Figure 17.) Some people think this is normal, desirable or not something the government should be expected to address. But if we do think that the fall in the wage share is a problem, then logically any solution must involve a period of time in which the wage share rises. A rising wage share involves, by definition, nominal wage increases greater than the sum of productivity growth and inflation. And if we take the view—supported by orthodox theory, empirical evidence and common sense—that nominal wage growth depends on the unemployment rate, this requires a period of unemployment below its full employment level. This is a critical point: It is impossible for the wage share to ever rise if the central bank will not allow a period of “excessive” wage growth. A rise in the wage share necessarily requires a period in which wages rise faster than would be consistent with long-term macroeconomic stability.




FIGURE 17 Total compensation (wages plus supplements, plus the labor share of proprietor’s income) as a share of national income. *Source: BLS.*

It is true that, in the past year, there has been an acceleration of wage growth relative to productivity, resulting in an increase in the wage share. But it would take several more years of similar wage growth to reverse the decline in the wage share of the 2000s. If the Fed treats any wage growth above productivity growth as a sign that growth is happening too rapidly, then the fall in the wage share cannot be reversed, no matter what happens to labor supply or technology. Logically, you cannot believe at the same time that the rising share of capital

in national income is a problem, and that there is no space for expansionary policy once full employment is reached. This is especially serious since the unemployment target is missed so often from the other side. If there are periods of excessively high unemployment but no periods of excessively low unemployment, there will be a kind of ratchet effect where the labor share can only go down, never up.

This is not part of textbook macroeconomics—the shares of labor and capital are normally considered to be long-term, supply-determined outcomes. But in recent years, there has been increasing recognition among macroeconomists that nominal wage growth below the sum of inflation and productivity growth necessarily implies a falling wage share (Binder and Rodrigue 2016; Bivens 2015). Since wage growth is closely related to macroeconomic variables, especially unemployment, the declining wage share is not purely the result of supply-side factors like technology, trade and education, but also reflects the macroeconomic policy chosen by the Fed.

Conclusions

The question of how close the economy is to potential is one of the most important questions for macroeconomic policy today. We should not take for granted that a low headline unemployment rate means that the economy has returned to full employment and full utilization of its productive capacity. We must critically examine this claim, using the full range of evidence available. When this is done, we have suggested, it is much harder to sustain the view that the economy is close to potential than it appears at first glance. There is a strong case to be made that, a full decade after the financial crisis, the fundamental problem in our economy remains a lack of demand.

In this paper, we have approached the question of potential output from a number of directions. We have shown that the current deviation of GDP from its long-term trend is larger and more sustained than any previous deviation since the end of World War II. The gap between actual GDP and its pre-2007 trend has not closed since the bottom of the recession. The apparent closing of the output gap is the result of downward adjustments in estimates of potential output. The bulk of this gap was not predicted by forecasters before the recession—a fact that many economists see as *prima facie* evidence that it is the result of the recession and not of autonomous structural factors.

We have shown that the fall in labor force participation, while it was largely predicted, cannot be solely explained by demographic factors. Less than half the fall in participation is attributable to the aging population; the majority represents lower participation *within* age groups. And, the aging of the population is offset by rising education levels, which historically have been associated with a higher level of labor force participation. If we take education as well as age into account, demographics explain none of the fall in participation rates.

Lower labor force participation and the associated slowdown in employment growth explain about half the output gap; the other half is explained by slower growth in output per worker, or productivity. While textbook theory suggests that productivity growth is a supply-side phenomenon, dependent on autonomous technological change and not on demand, we suggest that there are a number of reasons to expect strong demand to be associated with more rapid productivity gains, and weak demand to be associated with slower gains. There are a number of reasons to believe that weak demand has played a role in the exceptionally slow productivity growth of the past decade. First, the productivity slowdown has been spread out over a wide range of industries, including retail trade, services, transportation, durable and nondurable manufacturing, and other sectors. While the high-tech sectors that drove productivity gains in the 1990s have seen the largest falls in productivity growth, they account for only a relatively small portion of the overall slowdown. The breadth of the

slowdown makes it hard to interpret in purely technological terms, or as simply the end of the exceptional productivity growth of the tech boom period. Second, the combination of slow productivity growth and slow employment growth is historically unusual; normally these two trends move in offsetting ways. A simultaneous decline in both is, historically, more typical of business-cycle recessions than of extended periods of slower growth.

Finally, the state of the economy as a whole is inconsistent with a story of tightening supply constraints. Macroeconomic theory—as well as history and common sense—strongly suggests an autonomous decline in the available labor force or that the pace of technological progress should be associated with *higher* inflation. The signature of a negative supply shock is lower output combined with higher prices. Normally, a central bank responds to a negative supply shock with contractionary policy—higher interest rates—in order to bring the level of spending in the economy down to the new, lower level of productive capacity. Yet the past decade has seen consistently low inflation even in the face of ultra-expansionary monetary policy—exactly the opposite of what we would expect from a fall in aggregate supply.

Given this picture, we think there is a strong case that slow growth is largely, if not entirely, a product of weak aggregate demand. This implies that policies to increase spending—whether lower interest rates to encourage private borrowing, or improved incentives for business investment, or direct public spending, or measures to boost exports—could call forth greater employment and output. And it implies that supply-side measures to lower business costs might, in the absence of simultaneous steps to raise spending, move the economy toward weaker demand and deflation. It's also important, we argue, to take into account the relative risks of erring on each side. Higher inflation is, at worst, a problem we do not yet have and can easily deal with if and when it arises. Depressed output and employment, meanwhile, are serious and immediate problems. Deflation remains a real danger that, while remote, cannot clearly be dealt with by using existing tools. And one of the most important results of “overheating” is wages rising relative to productivity. But if real wage growth is never allowed to exceed real labor productivity growth, it is impossible for the wage share of national income to ever rise.

The idea that the economy has now reached supply constraints and is in danger of overheating should not become conventional wisdom without more critical scrutiny.

We acknowledge that much of the evidence we have presented here is suggestive, not dispositive. This paper is intended to be the start of a conversation, not the end of one. But it is vitally important that this conversation takes place. The idea that the economy has now

reached supply constraints and is in danger of overheating should not become conventional wisdom without more critical scrutiny.

This is true for several reasons. First, the Fed is a key policymaker that needs to be part of democratic debate. Second, if there is still a demand shortfall, policymakers need to be developing measures to boost spending. Opportunities may be fleeting, and if we believe the economy is already at potential, we may be lured into a false complacency about the continued need for bold proposals to encourage business investment, or to devote more public resources to health care, education and other pressing needs. Third, and most broadly, there is a danger in telling a dissatisfied public that this is the best we can do.

Historically, the Fed has been treated as an apolitical, technocratic, independent body whose decisions are not subject to political scrutiny. This was defensible in a period when its mission of maintaining full employment, stable prices and output near potential seemed to be consistently achieved, using a limited set of tools, and when fiscal policy offered a viable alternative in cases in which the Fed fell short. But “the Great Moderation” is clearly over, if not illusory. Even if one rejects the claim that there is a substantial output gap today, it is clear that output and employment have fallen dramatically short of potential over the past decade. The introduction of new tools of monetary policy, even if justified on pragmatic grounds, requires a higher level of democratic oversight; and, perhaps most critically, the evident incapacity of the other branches to carry out demand management necessarily shifts the focus of macroeconomic debate to the Fed. It might once have been acceptable to regard the level of potential output as a technical question on which citizens need not have a view; but if it ever was, that time has passed.

This is especially true given the priority the Fed evidently places on holding down inflation, as opposed to its other mandates. Rather than the “opportunistic disinflation” central banks have pursued for the past generation, there is a case for a policy of “cautious expansionism,” in which growth is pushed forward until clear signs of overheating are visible. More broadly, existing notions of central bank independence are not consistent with democratic norms, especially as the central banks’ role has expanded (Stiglitz 1998).

More broadly, the question of potential output conditions a great many other policy proposals. Any proposal to spend public money, or encourage the spending of private money, has an added benefit when output is below potential; and proposals that reduce public spending or discourage private spending will be more attractive when the economy is at or above potential. Extending health coverage to more Americans, investing in infrastructure and green energy, or improving access to higher education are ideas whose merits are independent of the state of the macro economy, but their associated price tags have a very different valence when spending more money is necessary to boost employment and output.

In the world of practical politics, windows for enacting such policies are likely to be narrow. If we believe that the economy requires someone to spend more money, we should have spending-boosting proposals at the ready. If we believe the economy is already at potential, we may be lured into a false sense of complacency about the continued need for bold proposals to encourage business investment, or to devote more public resources to health care, education and other pressing needs.

References

- Aaronson, Stephanie, Tomaz Cajner, Bruce Fallick, Felix Galbis-Reig, Christopher Smith, and William Wascher. 2014. "Labor force participation: Recent developments and future prospects." *Brookings Papers on Economic Activity* 2014, no. 2 (2014): 197-275. <https://muse.jhu.edu/article/577366/pdf> (April 20, 2017).
- Acharya, Sushant, Julien Bengui, Keshav Dogra and Shu Lin Wee. 2016. "Escaping Unemployment Traps." New York, FRB. <http://economics.rutgers.edu/downloads-hidden-menu/news-and-events/workshops/macroeconomic-theory/1360-keshavdogra/file> (April 20, 2017).
- Akin, S. Nuray, Ayhan Kose and Franziska Ohnsorge. 2014. "Estimation and Comparison of Output Gaps." Under Revision for *The World Bank Economic Review*.
- Alexander, Samuel. 2012. "The optimal material threshold: Toward an economics of sufficiency." *Real-World Economics Review* 61, no. 2. http://www.paecon.net/PAERReview/issue61/Alexander1_61.pdf (April 20, 2017).
- Armstrong, Philip, Andrew Glyn, and John Harrison. 1991. "Capitalism since 1945." Oxford: Basil Blackwell.
- Ball, Laurence M. 2009. "Hysteresis in unemployment: Old and new evidence." Cambridge, MA: National Bureau of Economic Research. <http://www.nber.org/papers/w14818> (April 20, 2017).
- Ball, Laurence M. 2014. "Long-term damage from the Great Recession in OECD countries." Cambridge, MA: National Bureau of Economic Research. <http://www.nber.org/papers/w20185> (April 20, 2017).
- Ball, Laurence. 2015. "Monetary Policy for a High-Pressure Economy." Washington, D.C.: Center for Budget and Policy Priorities. <http://www.cbpp.org/research/full-employment/monetary-policy-for-a-high-pressure-economy> (April 20, 2017).
- Ball, Laurence, Brad DeLong, and Larry Summers. 2014. "Fiscal policy and full employment." Center on Budget and Policy Priorities. http://larrysummers.com/wp-content/uploads/2014/04/delong_summers_ball.pdf (April 20, 2017).
- Binder, Carola, and Alex Rodrigue. 2016. "Monetary Rules and Targets: Finding the Best Path to Full Employment." Center on Budget and Policy Priorities. <http://www.cbpp.org/sites/default/files/atoms/files/9-1-16fullemployment.pdf> (April 20, 2017).
- Bivens, Josh. 2015. "A Vital Dashboard Indicator for Monetary Policy: Nominal Wage Targets." Center on Budget and Policy Priorities. http://www.cbpp.org/sites/default/files/atoms/files/6-15-15_full_emp-bivens.pdf (April 20, 2017).
- Blanchard, Olivier. 2016. "The U.S. Phillips Curve: Back to the 60s?" Washington, D.C.: Peterson Institute for International Economics. Policy Brief 1. <https://piie.com/publications/pb/pb16-1.pdf> (April 20, 2017).
- Blanchard, Olivier, Eugenio Cerutti, and Lawrence Summers. 2015. "Inflation and activity." In Unpublished manuscript, presented at the 2015 ECB Forum on Central Banking, Sintra, Portugal. https://2015.ecbforum.eu/up/ficheiros-bin2_ficheiro_en_0225454001432283160-368.pdf (April 20, 2017).
- Center for Economic and Policy Research. 2016. "March CPS Data." <http://ceprdata.org/cps-uniform-data-extracts/march-cps-supplement/march-cps-data/> (May 26, 2017).
- Congressional Budget Office. 2001. "CBO's Method for Estimating Potential Output: An Update." Washington, D.C.: Congressional Budget Office. <https://www.cbo.gov/sites/default/files/107th-congress-2001-2002/reports/potentialoutput.pdf> (April 20, 2017).

- Congressional Budget Office. 2007. "The Budget and Economic Outlook: Fiscal Years 2008 to 2017." Washington, D.C.: Congressional Budget Office. <https://www.cbo.gov/sites/default/files/110th-congress-2007-2008/reports/01-24-budgetoutlook.pdf> (April 20, 2017).
- Council of Economic Advisors. 2014. "The Labor Force Participation Rate since 2007." Washington, D.C.: Executive Office of the President of the United States. http://scholar.harvard.edu/files/stock/files/labor_force_participation.pdf (April 20, 2017).
- Cowie, Jefferson. 2010. *Stayin' Alive: The 1970s and the Last Days of the Working Class*. The New Press.
- DeLong, J. Bradford and Lawrence H. Summers. 2012. "Fiscal policy in a depressed economy." *Brookings Papers on Economic Activity*, (1), pp.233-297. <https://muse.jhu.edu/article/486372/summary> (April 20, 2017).
- Eatwell, John. 1997. "Effective Demand and Disguised Unemployment." In Michie J. and Grieve Smith J., eds, *Employment and Economic Performance*, pp.76-94. Oxford University Press.
- Engler, Philipp and Juha Tervala. 2016. "Hysteresis and Fiscal Policy." No. 1631. DIW Berlin, German Institute for Economic Research. http://www.diw.de/documents/publikationen/73/diw_01.c.549430.de/dp1631.pdf (April 20, 2017).
- Fatas, Antonio and Lawrence H. Summers. 2016. "Permanent Effects of Fiscal Consolidations." National Bureau of Economic Research Working Paper #22374. <http://www.nber.org/papers/w22374.pdf> (April 20, 2017).
- Fernald, John and Bing Wang. 2015. "The recent rise and fall of rapid productivity growth." *FRBSF Economic Letter* 9. <http://www.frbsf.org/economic-research/publications/economic-letter/2015/february/economic-growth-information-technology-factor-productivity/> (April 20, 2017).
- Furman, Jason. 2015. "Business investment in the United States: Facts, explanations, puzzles, and policies." Remarks by Jason Furman, Chairman Council of Economic Advisors. Progressive Policy Institute. https://obamawhitehouse.archives.gov/sites/default/files/page/files/20150930_business_investment_in_the_united_states.pdf. (April 20, 2017).
- 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 E. 2016. "Why Has the Unemployment Rate Fared Better than GDP Growth?" National Bureau of Economic Research. <http://web.stanford.edu/~rehall/Boston%20Fed%20paper%20June%202016.pdf> (April 20, 2017).
- Haltmaier, Jane. 2012. "Do Recessions Affect Potential Output?" Board of Governors of the Federal Reserve System International Finance Discussion Papers, No. 1066. <https://www.federalreserve.gov/pubs/ifdp/2012/1066/ifdp1066.pdf> (April 20, 2017).
- International Monetary Fund. 2017. "World Economic Outlook: Gaining Momentum?" World Economic and Financial Surveys. <https://www.imf.org/~media/Files/Publications/WEO/2017/April/pdf/text.ashx> (June 27, 2017).
- Kienzler, Daniel and Kai D. Schmid. 2014. "Hysteresis in Potential Output and Monetary Policy." *Scottish Journal of Political Economy*, 61(4): pp.371-396. <http://onlinelibrary.wiley.com/doi/10.1111/sjpe.12050/full> (April 20, 2017).
- Konczal, Mike and Marshall Steinbaum. 2016. "Declining Entrepreneurship, Labor Mobility, and Business Dynamism: A Demand-Side Approach". Roosevelt Institute Working Paper.

- Krugman, Paul. 2014. "Inflation Targets Reconsidered." Draft paper for ECB Sintra Conference. <https://online.ercep.com/media/attachments/inflation-targets-reconsidered-en-70089.pdf> (April 20, 2017).
- Lee, Timothy. 2017. "Why America needs more Teslas and fewer Ubers." Vox. <http://www.vox.com/new-money/2017/3/16/14939624/tesla-billion-raise-uber> (April 20, 2017).
- Mathy, Gabriel. 2015. "Hysteresis and Persistent Long-term Unemployment: Lessons from the Great Depression and World War 2." SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2574850 (April 20, 2017).
- The National Bureau of Economic Research. 2016. "NBER CPS Supplements." <http://www.nber.org/data/current-population-survey-data.html> (December 16, 2016).
- Pacitti, Aaron. 2017. "The Cost of Job Loss, Long-Term Unemployment, and Wage Growth." Unpublished working paper. <https://www.aeaweb.org/conference/2017/preliminary/paper/KhFtdG5s> (April 20, 2017).
- Pacitti, Aaron and Melissa Fichera. 2015. "The Ongoing Labor Market Depression." Challenge 58, no. 4.
- Roth, Jean. 2017. "Reading Current Population Survey (CPS) Data with SAS, SPSS, or Stata." May 11. http://www.nber.org/data/cps_progs.html (December 16, 2016).
- Stiglitz, Joseph. 1998. "Central Banking in a Democratic Society." De Economist. 146(2), pp. 199-226. <https://link.springer.com/article/10.1023%2FA%3A1003272907007?LI=true> (April 20, 2017).
- Stiglitz, Joseph E., Amartya Sen, and Jean-Paul Fitoussi. 2010. "Report by the commission on the measurement of economic performance and social progress." Paris: Commission on the Measurement of economic performance and social progress. http://library.bsl.org.au/jspui/bitstream/1/1267/1/Measurement_of_economic_performance_and_social_progress.pdf (April 20, 2017).
- Wilkinson, R.G. and Pickett, K. 2009. *The spirit level: Why more equal societies almost always do better*. London: Allen Lane.