

How Should We Think About this Recovery?

Jay Shambaugh
Member, Council of Economic Advisers

Macroeconomic Advisers' 26th Annual Policy Seminar
Washington, D.C.
September 14, 2016


This is an expanded version of these remarks as prepared for delivery.

Writing economic history in real time is hard. Data are sure to be revised, possibly altering conclusions soon after they are formed. As business cycles continue, our impressions of them may change. But, at the same time, as this expansion is now seven years old and as the Obama Administration is coming toward a conclusion in four months, it is natural to try to assess what has happened. Also, a proper appraisal matters more than just to history or to evaluations of past policy. Understanding where we are is crucial to our view of what policies are appropriate now and how we should forecast the next few years.

The Labor Market

There are many ways to look at a business cycle. We could, for example look first at the unemployment rate.

When this recovery began in July 2009, the unemployment rate was 9.5 percent and still rising. It would peak at 10.0 percent in October 2009, one of the many measures signaling this recession as the worst since the Great Depression. The rapid fall in the unemployment rate from this peak to 4.9 percent today has been truly remarkable. The unemployment rate fell faster than we collectively had imagined possible (see figure 1). In March of 2010, consensus forecasts suggested that the unemployment rate would be 6.5 percent by mid-2015; instead it was 5.3 percent. By March 2011, that projection had come down to 6.3, but that ended up being too pessimistic as well.



Time and again, the unemployment rate fell faster than we thought it would; each new forecast taking on board the new level but assuming the pace of decline would not be maintained. Now at 4.9 percent, the unemployment rate is well below where many set the steady state (or non-increasing inflation) unemployment rate a few years ago, and most assume it will continue to fall. The decline of the unemployment rate from its peak to today, an average rate of 0.75 percentage point per year, over 7+ years, is matched only by the expansion in the mid-1980s in terms of both pace and longevity. This drop in the unemployment rate suggests that by this metric, the recovery has exceeded economists' expectations.

Of course, the unemployment rate could have fallen for the wrong reasons. People could have simply left the labor market, and undoubtedly, early on, that was happening. But, over time, as the unemployment rate has been at a low level for a year or more, the estimated cyclical gap of the labor force participation rate appears to be closing. That is, the degree to which the recession is weighing on participation has been fading. The question is whether the participation rate has moved from the 2007Q4 business cycle peak to now in a way that suggests this cycle has been worse than others.

A number of studies have predicted the labor force participation rate (LFPR) or tried to estimate the impact of the business cycle on the rate. In 2006, Aaronson et al. (2006) predicted LFPR would be 62.5 percent in 2015. The annual average for 2015 was 62.6. In a 2014 update of that paper, Aaronson et al. predicted the 2016 labor force participation rate would be 62.7 percent. It is 62.8 today. In this sense, the developments in LFPR have been exactly what was expected even in the absence of a recession. These predictions were based on a combination of aging, other demographic changes, and pre-existing trends. This is of course just one of many forecasts of LFPR, and some of the drop in the participation rate during the recession was almost certainly due to the depth of the initial recession.

The Council of Economic Advisers built a model (CEA 2014) to decompose the decline in participation rate into aging trends, a cyclical component, and a residual. As the unemployment rate has come down, the cyclical effects have been fading, and as the LFPR has come up over the

past four quarters, the gap left from the aging trend has been shrinking as well. At this point in the recovery, the change in LFPR looks more and more like a typical cycle, adjusting for aging. There may be some residual greater than prior cycles that one would hope or expect to close, but one cannot write off the unemployment rate decline as strictly a function of labor force shifts.

Again, this is not to say there are no concerns with labor force participation. Far from it. The LFPR of prime age men has been on a steady decline for over 50 years. Decompositions like CEA's attribute some decline to aging, some more to cyclical weakness (which has largely faded), but a sizable residual remains. The residual decline in LFPR or the long-term trend in prime-age male participation are not new to this era, occurring fairly steadily over the past fifty years, but they are by no means salutatory. It would be preferable if LFPR did not decline more than simple demographics suggest it will. In fact, it would be preferable if prime age LFPR was increasing such that overall LFPR did not decline with demographics but was flat or even increased. Since late 2013, LFPR has been flat overall – suggesting a considerable cyclical rebound offset by demographic shifts.

Thus, we cannot simply attribute the unusual progress in the decline in the unemployment rate to shifts in participation. Even since October of 2013 when LFPR hit 62.8, the unemployment rate has continued to fall from 7.2 percent to 4.9 percent today, or a pace of 0.8 percentage point a year, just as fast as the overall unemployment rate recovery.

This is not to say that there is no slack in the labor market. The share of workers working part-time for economic reasons has fallen steadily throughout the recovery but is still elevated (3.8 percent in August vs. its average of 3.0 percent in the previous expansion) and as a result U-6 – the broadest measure of labor underutilization – is at 9.7 percent, slightly above the average of 9.1 percent previous expansion. So, there may be room to grow above trend. In addition, there is considerable uncertainty about the value of the NAIRU or the sensitivity of wages and prices to the gap between the unemployment rate and the NAIRU. (See CEA 2016).

Looking at employment tells a similar story of a robust recovery. Average job growth since the start of the labor market recovery has been 191,000 a month. We have witnessed the longest streak of positive total nonfarm job creation on record. Since the unemployment rate reached its pre-recession average of 5.3 percent in June 2015, job growth has averaged 207,000 a month, and has averaged 182,000 jobs per month so far in 2016.

In many ways, this most recent stretch of job growth has been surprising. CEA has estimated that for the unemployment rate to stay steady and LFPR to move with demographics would require net job creation of roughly 80,000 jobs a month this year. If the unemployment rate had kept moving down to 4.5 percent, it would require roughly 140,000 jobs a month. Instead, the unemployment rate has ticked down just one-tenth of a percentage point so far this year, and yet job creation has been far above these figures. The difference is primarily the uptick in the labor force participation rate by roughly 0.3 percentage point during the first three quarters of this year.

In some sense, job growth so far in 2016 seems like it has been at the outer edge of what one could have expected, but the job growth is coming primarily from pulling people back into the

labor force. Moreover, the pace has not been appreciably slowing, at least on a multi-month basis. In this sense, there does still appear to be some slack in the labor market. While prime-age LFPR has been trending down for decades, it could stop doing so, and if it did, there could be ample room to grow.

In the last few years, wages have been growing more quickly as well. While more work remains to raise the nominal growth rate, because inflation has been so low, real wages have been growing relatively quickly. This lower headline inflation is largely due to the substantial fall in oil prices that began in 2014. Some of the decline in oil prices is due to activity in the United States, where oil production rebounded from about 5 million barrels per day in 2008 to a peak of 9.6 million barrels per day in April 2015, largely attributable to technological innovations. This unexpected development has made the United States the world leader in oil production. Some is also likely due to the surprising slowdown in oil consumption in the United States. But certainly, oil is a global commodity, and part of the low price of oil – much like the lower commodity prices overall recently – is attributed to a slowdown in global growth, especially in some emerging markets. Core inflation has not been as notably slower this recovery compared with other recoveries, but headline inflation is the relevant measure for real wage growth. The question will be if as inflation rises towards the Federal Reserve's 2 percent PCE inflation target, will wages rise along with prices to keep this momentum in real wage growth.

An example of both recent progress on incomes as well as how quickly our view of a business cycle can change comes from the recently released data on income, poverty, and health. Real median household income rose 5.2 percent in 2015, the largest yearly growth on record, as household incomes rose at each reported decile for the first time in about a decade. These data combine recent wage gains and job gains and show real progress for all types of households, with the largest gains among households at the bottom of the income distribution.

Output Growth and Demographics

But jobs and the unemployment rate are not the only way we discuss the economy. The typical summary statistic is GDP.


Characterizing a recovery in real time with GDP is a particular challenge. While we can speak fairly confidently about what happened to unemployment rates in real time, and with a fairly short lag can say what has happened to employment levels, estimates of GDP can take a considerable amount of time to settle down. Some source data come with a lag on the order of years, not weeks, making firm pronouncements about how the economy is doing in a given quarter or even year challenging in real time. In addition, there is an increasing amount of debate about potential biases in the measurement of real GDP and while I would not be confident in assessing a large increase in any such bias, there is still increasing uncertainty as the economy becomes harder to measure.¹

¹ There is a debate as to whether measurement error is artificially lowering GDP and productivity growth statistics. These arguments tend to center around developments in high-tech products or services where quality adjustments may be more difficult to measure. In this talk, I proceed from the assumption that there has not been a substantial

Nevertheless, looking at real GDP growth over the full sweep of a business cycle gives a fairly good sense of the pace of economic progress over that period. There are many things that are important in life beyond GDP, but in the end, GDP gives us a measure of the output of an economy, and gives us a sense of the growth in the economic resources a country has at its disposal.

At this point in the current cycle, we have seen a clear recovery in output. Real GDP was growing again by the third quarter of 2009, and by the second quarter of this year, real GDP was 10.5 percent higher than the pre-recession peak. Real GDP per capita was above its pre-crisis peak within 54 months from the end of the recession. In contrast, the euro area still cannot make that claim. The U.S. recovery has also outpaced that of Japan.

Figure 2: Real GDP per Capita - Euro Area, United States, and Japan



Source: BEA (U.S.), Economic & Social Research Institute (Japan), Eurostat and UN Population Statistics (Euro Area); Haver and CEA calculations.

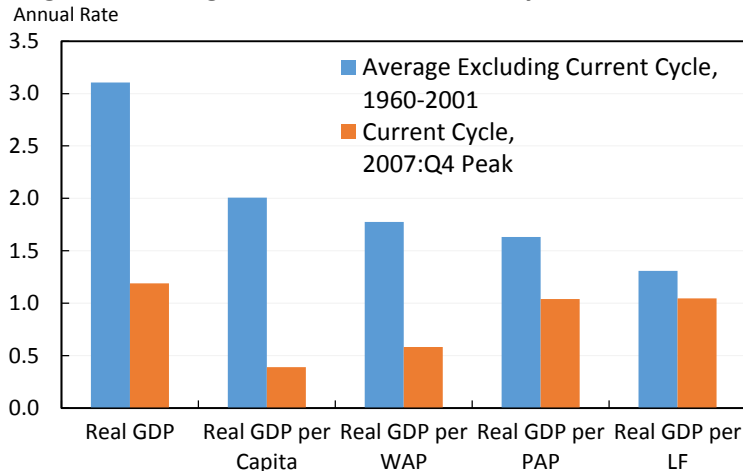
When looking at this business cycle, a number of commentators have focused on the pace of the recovery relative to other business cycles. In some ways, comparing business cycles peak to peak has a compelling logic to those of us who study macroeconomic history. If a crash was very harsh, and the rebound fast, looking peak to peak will balance those features. In a world where the macroeconomy oscillates around a long run trend with boom times where output is above its potential real GDP followed by weak episodes where it is below a potential growth trend line by an offsetting amount, such a procedure makes perfect sense. But it certainly leaves out the question of what caused the shock, whether the prior peak had been fueled by speculation or asset bubbles, and most of all, whether structural changes may have affected the path of the economy.

With respect to real GDP growth, a measure that I will argue is misleading if no adjustments are done, the current recovery appears not to have been strong. Because a number of crucial features

increase in that mismeasurement in this business cycle that would change growth in this business cycle relative to others.

vary across cycles, these raw GDP comparisons are inappropriate. The one I'd like to focus on in these remarks is the role of demographics. As we take into account a number of demographic features, we see the recent recovery looks far more typical.


Figure 3: Average and Current Business-Cycle Peak to Peak



Note: Working age population (WAP) is aged 15-64; prime age population (PAP) is aged 25-54; labor force (LF). Peak start dates: 1960q2, 1969q4, 1973q4, 1980q1, 1981q3, 1990q3, 2001q1.
Source: BEA, NBER, Census, BLS, UN Population Statistics; Haver and CEA calculations.

Figure 3 shows peak to peak comparisons, but trough to peak comparisons may also be interesting (see figure 4). Using peak to trough calculations has its own difficulties: if the economy often bounces out of a recession and then slows before a cycle ends, then a shorter recovery, with a quick burst of growth followed by a recession quickly may wind up with what appears to be a rapid growth cycle, but this gives us a sense of the pace of the expansion itself, not just the overall cycle (which includes the downturn). The chart below shows a similar comparison across different growth measures, but calculated from business-cycle trough to peak. As some observers have pointed out, using either of these measures of simple GDP growth, the current recovery appears slower than the average of economic history. But these measures miss vital context and are not the right way to look at this recovery without adding more considerations.

Figure 4: Average and Current Business-Cycle Trough to Peak



Note: Working age population (WAP) is aged 15-64; prime age population (PAP) is aged 25-54; labor force (LF). Trough start dates: 1961q1, 1970q4, 1975q1, 1980q3, 1982q4, 1991q1, 2009:Q2. Source: BEA, NBER, Census, BLS, UN Population Statistics; Haver and CEA calculations.

To factor in demographics, the first adjustment we would want to make is to look at growth in per-capita terms. This simple adjustment makes a sizable difference to how we view different business cycles. Slower population growth over this business cycle makes comparisons of raw GDP growth in the last 7 years to business cycles in the 1960s inappropriate. The average population growth rate across previous business cycles in the last half century was 1.1 percent compared with 0.8 percent so far in this business cycle. Comparing business cycles in terms of real GDP per capita meaningfully narrows the gap in growth rates.

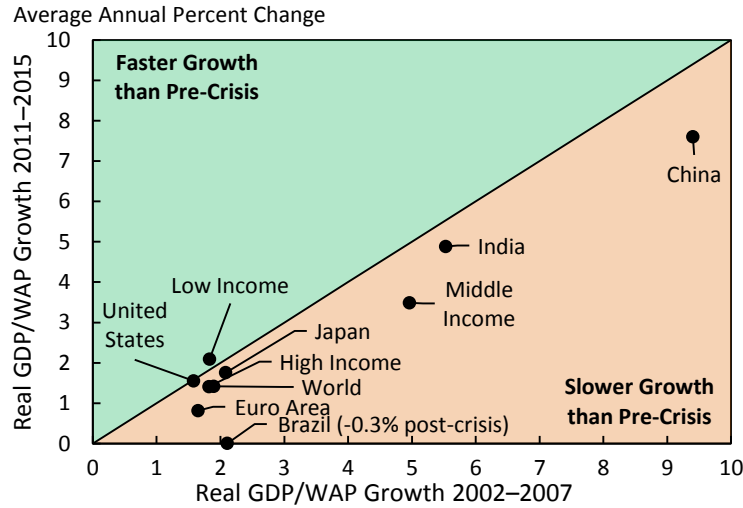
But, the change in the current population growth rate is in many ways less important than the changes that have taken place in the demographic distribution of the current U.S. population. The crucial question for output is not how many people are in the society, but how many people are working. The working age population (ages 16-64) grew 1.4 percent at an annual rate in the 1960s through 1980's, but just 0.6 percent so far during this recovery. Thus, we could consider the growth rate of GDP per working age person in the economy as a way to control for such variations across cycles.

On this measure, the gap from the current recovery to history closes even more quickly. By this measure, GDP has been growing peak to peak at 0.6 percent, and since the recovery began at 1.5 percent. This is much closer to the average of the previous cycles compared with looking at unadjusted GDP and is roughly the same pace as the recovery following the 2001 recession.

That the current recovery looks somewhat like the early to mid-2000s on with regards to the growth of GDP per working-age person measure is rare when looking across countries. Nearly every country or group of countries has had a lower growth rate of GDP per working age person in the 5 years after the bust and boom of the crisis had ended (2011-2015) compared with the six years preceding the bust (2002-2007). As noted, there are many reasons we might expect different growth rates in these eras, including the fact that countries had room to grow quickly given the still present output gaps in 2011 following the Great Recession. But, when looking across countries, after taking into account demographics, many countries experienced a sizable slowdown. All these countries were faced with the after effects of the financial crisis and all

these countries had room to grow to make up ground lost in the crisis, but nearly all wound up growing slower – adjusting for changes in the working age population – after the crisis than they did before. The United States is one exception. Lower income countries are another.

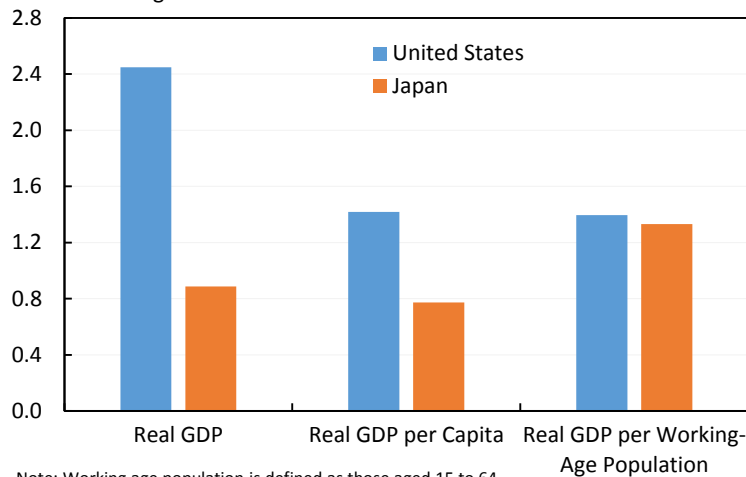
Figure 5: Real GDP per Working Age Population, Pre- vs Post-Crisis



Taking the same set of steps to control for demographics makes Japan’s growth after its economic slowdown in the early 1990s look much better. Recently, headline economic growth in Japan has been stagnant (and forecasted to remain so) in large part because growth in the working-age population has lagged behind growth in the total population for the last 20 years (Mühlesisen and Faruqee 2001). The U.S. economy grew more than twice as fast as Japan’s from 1991 to 2015. However, simply controlling for population by comparing growth of real GDP per capita leads to more similar growth rates (1.4 percent versus 0.8 percent). Even more striking is that when examining GDP per working-age person, the growth rates are essentially the same. This highlights that even if a country is doing reasonably well conditional on its demographics—as Japan has—it still means slow growth over time if too few workers enter the labor force. And even if income per capita is rising, slow overall growth due to slow population growth can greatly increase the challenges associated with government debt and financing future government commitments (Karam, Muir, Pereira, and Tuladhar 2011).

Figure 6: Real GDP Growth and Demographic Trends, 1991-2015

25-Year Average Annual Growth Rate



Note: Working age population is defined as those aged 15 to 64.


Source: U.S. Bureau of Economic Analysis; Japan Cabinet Office; World Bank; CEA calculation.

Even looking at working age population (16-64) may be an insufficient adjustment. Participation rates are much lower in young age groups where many people are in school and in older groups where people have retired or faced physical difficulties removing them from the labor market. Thus, one could look at GDP per prime age worker (25-54). Changes in the number of workers in those older and younger cohorts would still be altering GDP, so this is not a perfect measure either, but it too helps remove some of the impact of demographic shifts when comparing across eras.

The final step that is useful when looking across U.S. recoveries when thinking about demographics is to consider broader shifts in the labor force. For over half a century, the prime-age male labor force participation rate has been falling in the United States (see CEA 2016 for extensive discussion). But for many years, this was offset and obscured by the rapid increase in women's labor force participation from below 40 percent in 1950 to above 60 percent by the mid-1990s. The combination meant that the labor force was growing even conditional on the working-age population growth rate, that is, the overall working age LFPR was growing. Thus, recoveries in the 1950s and 1960s, and in particular in the 1970s and 1980s were in part driven by the rapid shift of women into the labor force.


As figure 7 shows, the rapid increase in the prime-age participation rate slowed in the 1990s and peaked at 84.1 percent in 1997. Since then, prime age labor force participation has declined to 80.9 percent in 2015 (though it is up to 81.3 percent as of August of this year).

Figure 7: Labor Force Participation Rate Aged 25-54, 1950 to 2015
Share of Population, Aged 25-54



If we take into account the number of people in the labor force, not just the working-age population, the growth rates across cycles look much more similar. The historical average is 1.3 percent compared with 1.0 percent in the most recent period. From trough to peak the figures are 2.3 percent and 1.9 percent respectively. As figure 8 shows, on a trough to peak basis, the current recovery looks quite similar to nearly every other recovery with the exceptions of the 1960s and 1980s expansions.²

Figure 8: Real GDP per Person in the Labor Force, Trough to Peak
Index, Business-Cycle Trough = 100



² This analysis reaches a similar conclusion to Stock and Watson (2012) who find that changes in potential (largely due to changes in demographics) can explain roughly three-quarters of the difference in growth early in this recovery compared with other recoveries.

To be clear, the Administration has continually pushed for policies to strengthen economic growth (only some of which have been approved by the Congress), and the recovery could have been faster if all of these policies had been approved. At a number of points in time, the Administration tried to provide more support to the economy – something that Ben Bernanke was explicitly asking Congress to do – but were rebuffed by Congress. A great deal of stimulus was put into the economy, not just the Recovery Act of 2009, but the payroll tax cuts, unemployment insurance extensions, aid to States, and extensions of tax cuts and tax credits as well. But the Administration had wanted to do more. At almost every turn, the Administration has suggested increased infrastructure spending – in late 2009, in each Budget, in the American Jobs Act proposal in 2011, continuing through the President's fiscal year 2017 Budget proposal. If one accepts that fiscal stimulus generates short-term growth – something I would wager most macroeconomists believe – then it is straightforward that had more of these proposals been enacted, the recovery would have been faster at various points in time, and had policies that also increase long run growth like infrastructure, business tax reform, or those that build human capital, the overall growth rate would likely have been higher as well.


Productivity Growth and Demographics

The very strong labor market recovery paired with a more modest GDP recovery adds up to weak labor productivity growth. Since the 2007 business-cycle peak, productivity has been growing at 1.0 percent a year, compared with 2.1 percent in the business cycles from 1960 to 2001. International evidence is again useful. In 30 out of 31 advanced economies, productivity growth slowed in the period 2005 to 2015 compared with 1995 to 2005—and the United States has had the fastest productivity growth of any of the G-7 economies in the last decade.

It is hard to put a uniquely country specific cause at the center of this decline in productivity when it appears internationalist be global in nature. Two common threads are the fact that demographic shifts are widespread across countries and investment has been weak globally.


The productivity slowdown in the United States has not been primarily a collapse in innovation or new ideas. In fact, more than half of the slowdown comes from a simple lack of capital deepening. And this explanation is common across many major advanced economies. The slowdown in capital deepening from its pace in 1994-2004 to 2004-2014 was in fact larger in Germany, Japan, and the UK than it was in the United States.

Figure 9: Change in Growth in Components of Productivity in the G-7, 1994-2004 to 2004-2014



Also, though, the demographic challenge we face is common across many countries. For the last 40 years, the world has ridden a wave of a rising working-age population that has powered growth. In Europe, the working-age population is now declining – not just as a share of the labor force, but in raw numbers. In many other regions, that peak is coming soon or has already happened. (South Asia and Africa are two notable exceptions).

Figure 10: Actual and Forecasted Working Age Population by Region

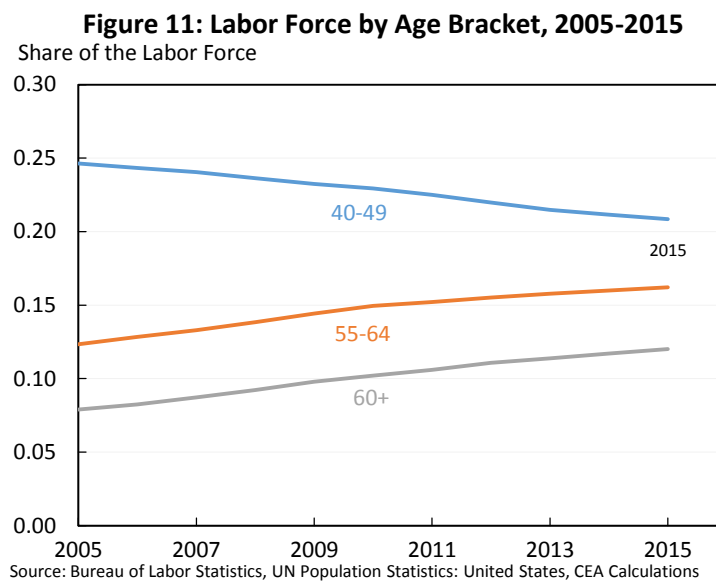


The connection of demographics to productivity growth is less obvious than its connection to labor force participation, but it may be a partial explanation for the productivity slowdown in the United States in the last decade. The idea that demography is correlated with things other than labor force participation is not new, financial analysts have long noted the correlation of the “Yuppie / Nerd” ratio (the ratio of 25-34 year olds to 40-54 year olds) and long-term interest

rates as people early in their working lives may be more apt to borrow to finance homes and those in their later working years may be saving more intensively.

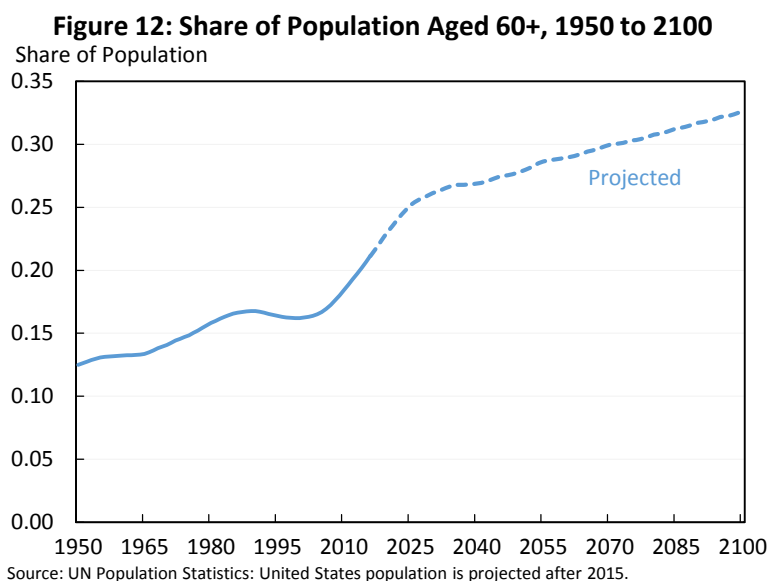
Over the last decade, economists have been seeing the link to productivity as well. James Feyrer argued in 2007 that demography was not just something that was changing the number of workers or the number of hours worked, but that demographic factors had an important role to play in productivity growth as well. His primary result was that, when considering the full age distribution of the labor force, in places where the proportion of 40-49 year-olds in the work force was growing relative to other cohorts, productivity growth was unusually fast. While this may seem like a statistical quirk (especially to those of you in the audience *not* in that demographic), Feyrer discussed in a series of papers possible explanations including the notion that the larger number of 40-49 year olds meant a larger number of qualified managers such that managerial talent was on average higher. Given the importance of management to the effective use of labor, this meant that productivity growth was higher in these countries or eras.

Feyrer’s work also had an interesting observation and disturbing prediction. Japan had faced a productivity slowdown following its crash, but not just due to a financial crisis, instead due in part to shifting demographics and an aging population leaving fewer 40-49 year olds as a share of the labor force. The disturbing prediction was that beginning in 2007, the United States would face an extended period where the number of 40-49 year olds would be shrinking as a share of the labor force, implying slower productivity growth. As figure 11 shows, the 40-49 cohort is shrinking over this past decade while the 55-64 and 60+ cohorts are rising. The coefficients in Feyrer’s paper multiplied by the shift in U.S. cohort shares implied that U.S. productivity growth would be 0.4 percentage point per year lower than otherwise during the 2005-15 interval. He also creates an instrument for the different demographic shares of the labor force, and using coefficients for that implies that demographics dragged down Total Factor Productivity by an average of 0.5 percentage point during the last 10 years.



A number of other papers have noted similar correlations of demographic patterns and labor productivity growth. In a recent blog post from the IMF, Shakhair Aiyar, Christian Ebeke, and Xiaobo Shao note an impending worry in Europe. They show that as the share of 55-64 year olds rise, labor productivity drops. Their results imply that if the share rises by 5 percentage points, total factor productivity will be cumulatively 2-4 percent lower at the end of that period. The note is meant to be a warning for many countries in Europe that are about to face a rising share of 55-64 year olds, but in the United States, it is in many ways more a story about the last decade than about the next. Since 2006, the share of 55-64 year olds in the United States rose by 3.9 percentage points, implying productivity 1.5-3.1 percent lower than it otherwise would have been, or roughly lower average productivity growth of 0.15 to 0.3 percentage point a year.

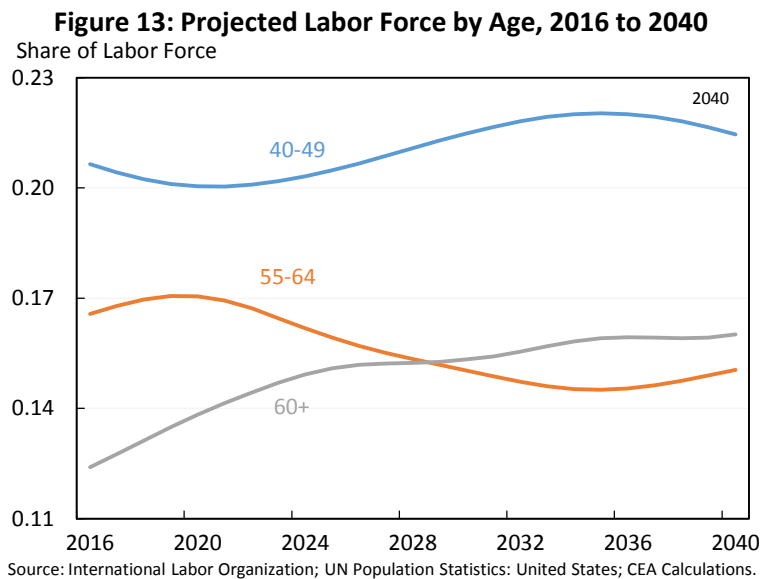
Finally, in another paper Maestas, Mullen, and Powell (2016) suggests the slowdown in total factor productivity (and economic growth) is related to the increase in the share of the population that is over 60. They distinguish between the impact due to labor productivity and the impact due to slower labor force growth. They find sizable impacts on GDP per capita when the population over 60 increases as a share of the total population. Some of this, of course, is due to shifts in labor force participation. But, surprisingly, they find that two-thirds of the effect is due to changes in output per hour worked. In total they find that a 10 percent increase in the share of the population over 60 reduces GDP per hour worked by a cumulative 3.4 percent. As the share of the U.S. population over 60 increased by 4.0 percentage points from 2005 to 2015 (see figure 12), or about 25 percent, this implies a cumulative 7.3-percent reduction in productivity growth over that period or a 0.8 percentage-point per year impact.



It is important to note that the effect of a falling 40-49 share are almost certainly not additive with the effects of a rising share for older workers. A smaller share of 40 year olds may have negative impacts in part because it implies a larger share of 55-64 year olds or *vice versa*.³

³ In addition, some results are for labor productivity, some for total factor productivity, and some may have concerns with endogeneity of the population cohorts given the temptation to move in response to economic potential. But the overall picture they present is one where the labor force participation impact of aging is not the only, or even largest, impact. In addition, while Maestas et al and Aiyar et al focus on just one segment of the population, Feyrer looks at

Rather than see these as something to be combined, they seem to point in the same direction. Demographics have been exerting a very real drag on U.S. labor productivity growth for the last decade. It seems unlikely demographic changes explain the entire 1.1 percentage points slower annual productivity growth rates in this era compared with the past, but they may explain a sizable portion. Interestingly, these results also suggest that the drag is nearly done. The share of 55-64 year olds in the labor force looks to be roughly stable going forward, as does the share of 40-49 year olds (see figure 13). The share of the population that is over 60 is still rising as is the share of 60+ in the labor force. Taken together, these projected demographic forces imply that some of the productivity slowdown over the past few years may abate as the demographic mix within the workforce stabilizes, though the share of 60+ would suggest the effect may continue.




Another important way in which demographics can affect a recovery is through the impact on the long-run expected growth rate. Japan's experience presents an important warning. Even though Japan's economy has grown at a respectable pace on a per-working-age person basis, and even though its real GDP per capita is relatively high, its low overall growth rate presents lower incentives to invest and makes the dynamics of stabilizing the debt-to-GDP ratio more challenging. Following the massive shift in demographics, many problems continue to plague the economy. Long-run fiscal solvency looks more difficult because of slower baseline growth. Investment and productivity have been low as well. Even if one can explain the lower growth on a per-person basis or explain the lower productivity growth, the lower expected GDP growth in the future will tend to depress investment as well given the degree to which investment is a function of output growth.

Investment in the United States rebounded quickly in this recovery, but after a few years slowed notably. Over the last year, business fixed investment has contracted. Rather than appeal to various arguments about finance or uncertainty or corporate choices, it is useful to begin from a

the full distribution of the population. Finally, the Maestes et al paper uses the share of the population not the labor force as an explanatory variable though the link to output per hour rather than output per capita is somewhat less intuitive in that case.

basic model of investment, the accelerator model. This model assumes that investment growth is a function of change in the growth of real GDP. (Because firms target a level of the capital stock that moves with the overall level of GDP). A simple specification of the accelerator model, augmented with global growth, generates a fairly close match to actual investment growth.⁴ Thus, we could suggest that the lower LFPR due to aging and lower productivity growth due to demographic shifts have put downward pressure on investment growth by lowering GDP growth as well.

Figure 14: Business Fixed Investment - Actual vs. Accelerator



Policy Implications

There are of course other important caveats when comparing eras. On the one hand, the current cycle began with a deeper recession, which should have allowed a stronger rebound in our calculations beginning from the cyclical trough. On the other hand, study after study has suggested that economies recover more slowly from financial crises and housing busts. Numerous papers – work from the IMF early on in the crisis, work by academics, work from the Federal Reserve⁵, have tried to determine what one should expect from recoveries after credit booms or financial crises. This rich and interesting literature is beyond the scope of my remarks today. I think it is safe to say that on net, this literature has suggested that recoveries from financial crises have unique challenges and may require a longer recovery period.

⁴ One can estimate the model using lags of the growth of output – where the changes in sign of the coefficients across lags will pick up the importance of the change in output growth – or estimate the model using the change in output growth as the explanatory variable. Both models fit the past twenty five years fairly well. We also include foreign growth aggregated using U.S. trade weights. The series used is the four quarter percentage change or the change in that value, for first and second differences specifications respectively. In addition, models shown also includes lags of changes in investment growth and an error correction term for both output and investment.

⁵ See for example IMF (2009); Cecchetti et al (2009); Jorda, Schularick, and Taylor (2013). Romer and Romer (2015) argue that financial crises do not necessarily trigger potential output losses and Martin, Munyan, and Wilson (2015) find that financial crises do not necessarily generate different types of recessions, though they find deeper recessions do take longer to recover from and can be aided by fiscal stimulus.

One reason I have not emphasized this strain of the literature is that while many studies do find a slowdown in growth following a financial crisis on average, not all do, and the reality is that many of those that find a slowdown on average also find a number of episodes where growth does not slow at all. Policy in the wake of a crisis is quite important and can meaningfully shift the course of a recovery, even when the housing and financial sectors prompted the crash.

And, looking at this recovery – in particular in the context of other countries around the globe – a great deal of evidence suggests that policy can in fact reorient an economy to a faster path following a financial crisis. In the United States there was immediate and bold fiscal, financial, and monetary policies that responded to the crisis. This helped start the rebound faster than one would have expected given the depth of the shock and likely saved millions of jobs and trillions of lost output. In other countries where either fiscal policy was constrained by markets or ideology, or where a monetary policy response was delayed, growth was slower. So, the average output response to a financial crisis really should not be viewed as destiny. Instead, the outcomes may depend importantly on policy.

If this recession's connection to a financial crisis was the only important factor reducing growth compared with the past, then fixing the financial system is what matters. But the United States has not had the experience of Japan in 1991 or Europe in 2011. The United States did tackle the banking system problems in a head on manner, aggressively recapitalizing banks, providing liquidity, and reforming financial rules. Interest rates are low and non-financial firms do in fact have both available liquid funds and the ability to borrow.

Instead, the demographic shock and its impact on productivity as well as labor force participation paired with weak global growth seem a better explanation than the financial crisis alone for slower headline growth than in the past. I would argue that as relevant as the financial crisis literature may be, taking demographics into account properly is if anything more important when comparing across eras.

This would suggest that policies that can lift productivity growth and labor force participation are crucial and that support for growth to lift output growth (and with it investment which would also help lift productivity growth and participation) is needed. Some of the more obvious ways to approach these issues would be to expand infrastructure spending in ways that boost productivity. Interest rates are low, infrastructure is degrading, and we need to update our transportation system in many parts of the country to handle changing populations as well as make them more climate friendly. Put together, these considerations all point to clear benefits from more infrastructure spending along the lines the President proposed in the 21st century infrastructure plan in the Fiscal Year 2017 Budget. Business tax reform could lift investment incentives. Making higher education more affordable would lift labor quality. Trade agreements could lift exports and productivity. Finally, the Administration has worked to lift global demand in forums like the G-20 to try to alleviate the drag of slower global growth on U.S. exports and investment.

At the same time, policies that try to catch workers when they fall are needed to stabilize labor force participation among prime-age workers. The research of Autor, Dorn, and Hanson (2013) that some have called the “China shock literature” has more to say about problems we see

keeping workers in the U.S. labor market after a negative shock than it really says about trade or any one country. Wage insurance, streamlining unemployment insurance, training initiatives, and policies to help communities that are particularly affected by a given shock could all help make sure people do not drop out when they face a bad shock. And more family-friendly policies could make it easier for parents to work, lifting labor force participation.

All of these policies would help address the longer-term structural challenges in the U.S. economy and ensure that our economy is more resilient when the next economic downturn strikes.

These longer-term trends have also affected the shape of this recovery. Productivity growth slowed in the mid-2000s. Demographics may be playing a role, but overall, economists do not fully understand the reasons behind slow growth over the past decade in output per hour. As I've already noted, some of the trends in labor force participation apart from demographics stretch back over fifty years. That is important when we try to think about their causes or ways we need to address them, but it does not change the impact on incomes for many individuals and families. Furthermore, trends in inequality over the last few decades have created challenges for middle-class families across both recessions and recoveries. Even with sizable shifts in the tax code that have lowered the tax liabilities of many in the bottom half of the distribution and raised them on those in the top 1 percent along with other meaningful policy changes over the past seven and a half years, the shift in inequality over decades was so large that much work remains.

These longstanding trends color the way we look at the economy today, but when we think about the macroeconomics of this cycle we should be clear, in many ways it has been a solid recovery.

Conclusion

So, on net, how should we judge this recovery? It is a recovery that was preceded by a global financial crisis on a scale not seen in generations. And yet, the unemployment rate fell rapidly from its 2009 peak, job growth has been strong, and real wages are growing. In many important ways, the recovery has exceeded the expectations of economists. Examining the rate of GDP growth on an apples-to-apples basis, taking into account demographic changes and long-term trends in labor force participation, also shows a recovery in line with nearly all previous expansions in the post-war period. These considerations not only help provide the most relevant economic comparison but also helps to explain some of the trends we have witnessed in productivity growth as well.

On the other hand, our economy continues to grow and in the coming years, our understanding of this business cycle and expansion will almost certainly change. And, as data on GDP, productivity, and other crucial measures revise over time, our picture of even these first seven years of the recovery may evolve as well.

That's what makes real time macro and forecasting so fascinating.

References

- Aaronson, Stephanie, Bruce Fallick, Andrew Figura, Jonathan Pingle, and William Wascher. 2006. "The Recent Decline in the Labor Force Participation Rate and Its Implications for
- 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* (Fall 2014): 197-275
- Cecchetti, Stephen, Marion Kohler, and Christian Upper. 2009. "Financial Crises and Economic Activity." *NBER Working Paper* No. Working Paper 15379.
- Council of Economic Advisers (CEA). 2014. "[The Labor Force Participation Rate Since 2007: Causes and Policy Implications.](#)"
- _____. 2016. "[The Long-Term Decline in Prime-Age Male Labor Force Participation.](#)" Report.
- David, H., David Dorn, and Gordon H. Hanson. 2013. "The China syndrome: Local labor market effects of import competition in the United States." *The American Economic Review* 103(6): 2121-2168.
- Eggertsson, Gauti B., Neil R. Mehrotra, Sanjay R. Singh, and Lawrence H. Summers. 2016. "[A Contagious Malady? Open Economy Dimensions of Secular Stagnation.](#)" NBER Working Paper No. 22299.
- Feyrer, James. 2007. "Demographics and Productivity." *Review of Economics and Statistics* 89 (1): 100-109.
- Karam, Philippe, Dirk Muir, Joana Pereira, and Anita Tuladhar. 2011 "Beyond Retirees." *Finance and Development* 48 (2): 12–15.
- Jordà, Òscar, Moritz Schularick, and Taylor Alan. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45 (S2): 3-28.
- Maestas, Nicole, Kathleen J. Mullen, and David Powell. 2016. "[The Effect of Population Aging on Economic Growth, the Labor Force, and Productivity.](#)" NBER Working Paper No. 22452.
- Martin, Robert, Teyanna Munyan, and Beth Anne Wilson. 2015. "Potential output and recessions: are we fooling ourselves?" *FRB International Finance Discussion Paper* 1145.
- Mühleisen, Martin and Hamid Faruqee. 2001 "Japan: Population Aging and the Fiscal Challenge." *Finance and Development* 38 (1).
- Romer, Christina D., and David H. Romer. 2015. "New evidence on the impact of financial crises in advanced countries." NBER Working Paper No. 21021.