

STEPHANIE AARONSON

Federal Reserve Board

TOMAZ CAJNER

Federal Reserve Board

BRUCE FALICK

Federal Reserve Bank of Cleveland

FELIX GALBIS-REIG

Federal Reserve Board

CHRISTOPHER SMITH

Federal Reserve Board

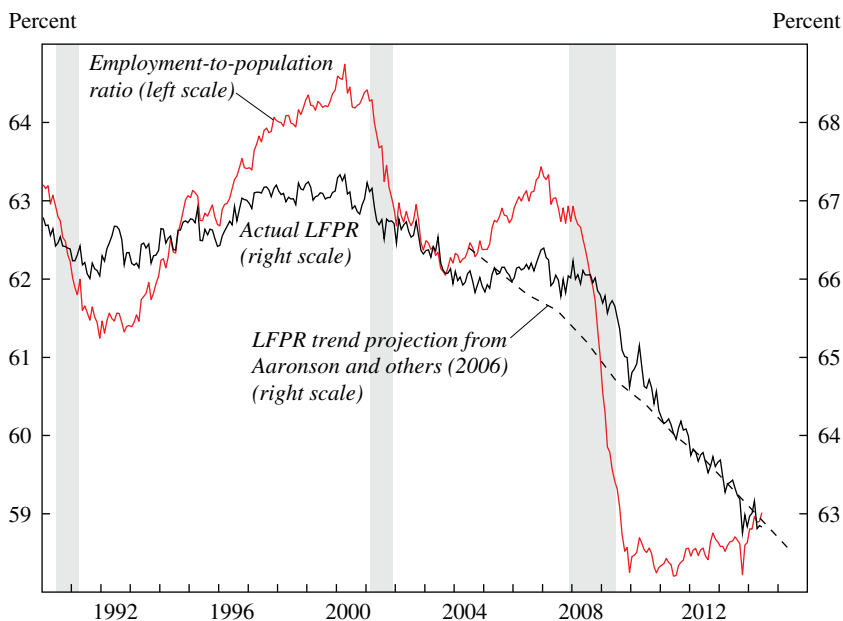
WILLIAM WASCHER

Federal Reserve Board

Labor Force Participation: Recent Developments and Future Prospects

ABSTRACT Since 2007, the labor force participation rate has fallen from about 66 percent to about 63 percent. The sources of this decline have been widely debated among academics and policymakers, with some arguing that the participation rate is depressed due to weak labor demand while others argue that the decline was inevitable due to structural forces such as the aging of the population. In this paper, we use a variety of approaches to assess reasons for the decline in participation. Although these approaches yield somewhat different estimates of the extent to which the recent decline in participation reflects cyclical weakness rather than structural factors, our overall assessment is that much of the decline is structural in nature. As a result, while we believe some of the participation rate's current low level is indicative of labor market slack, we do not expect the rate to substantially increase from current levels as labor market conditions continue to improve.

More than five years after the Great Recession ended, the labor market has, by many metrics, finally shown substantial improvement. As of mid-2014, the unemployment rate had fallen nearly 4 percentage points from the peak reached in late 2009, and the number of nonfarm payroll jobs had returned to prerecession levels. However, one lingering concern is the lack of recovery in the labor force participation rate and the concomitant absence of a significant rise in the percentage of the working-age

Figure 1. Labor Force Participation Rate, Actual and Predicted, 1990–2014

Source: For actual LFPR and employment-to-population ratio, Bureau of Labor Statistics; for LFPR projection, Aaronson and others (2006).

population that is employed. In particular, the labor force participation rate fell from about 66 percent of the population in 2007 to about 63 percent over the first half of 2014, while the employment-to-population ratio as of mid-2014 stood at 59 percent, only about $\frac{1}{2}$ percentage point above its low point in the wake of the recession (figure 1).

To an important extent, this decline in the labor force participation rate likely reflects the ongoing influence of the aging of the population, which was one focus of a Brookings paper written nearly a decade ago by several of the present authors (Aaronson and others 2006). Indeed, in that paper, we predicted further declines in the participation rate over the subsequent decade based on the aging of the population and longer-run trends in participation that, it seemed to us, were likely to hold down participation going forward. However, population aging cannot account for the entire decline in the aggregate participation rate, and the deep recession that was precipitated by the financial crisis, along with the slow economic recovery that has followed, have led some observers to ask whether cyclical factors have

played an important role as well and, if so, whether many individuals who dropped out of the labor force because they became discouraged about their job prospects may eventually re-enter the workforce as the labor market continues to strengthen.

The answers to these questions have important implications for government policies. If much of the decline in the participation rate can be reversed (or a further decline prevented) by a sufficiently tight labor market, arguably policymakers should take the low level of the participation rate into account in designing countercyclical policy actions. However, some of the decline in the participation rate might not be amenable to countercyclical policies. We will refer to this latter portion of the decline as “structural” in nature, and believe that these structural factors present a different set of challenges for policymakers. To the extent that these factors are caused by obstacles faced by individuals who would like to work or by disincentives to work, policymakers would be well advised to look for other ways to mitigate them. In addition, some of these structural factors may be unpreventable (such as aging of the population) or undesirable to reverse (such as higher school enrollment rates among the young).

Our primary aim in this paper is to assess explanations for the decline in the participation rate since the onset of the Great Recession, since disentangling the cyclical and structural changes over the past seven or so years is particularly complicated and, again, has important policy implications. However, because participation rates have actually been falling for some demographic groups since well before the recession began, at times our analysis necessarily extends to earlier periods in order to properly frame more recent developments.

We begin, in section I, by summarizing some of the alternative views of the recent decline in labor force participation, highlighting the wide range of often contradictory conclusions these studies have reached. In section II we take multiple approaches to assessing the recent decline. We first examine a number of explanations for the decline in participation that may be structural in nature. Foremost among them is the aging of the population, which *a priori* seems likely to have been a significant contributor. In fact, we find that it can explain nearly half of the decline in participation between the fourth quarter of 2007 and mid-2014. We also consider whether the pre-recession decline in participation within some demographic groups, such as younger adults and less-educated prime-age males, implies that participation rates for these groups would have fallen even in the absence of the recession, and we explore the importance of disability insurance take-up as

an additional structural factor. We find that these factors also contributed to the downward trend in participation over these years, so that structural factors could conceivably explain almost all of the decline in participation since 2007.

In section III we approach the question from a different angle and attempt to directly assess the cyclical component of the recent decline in participation, using evidence from a series of cross-state panel regressions. Whereas our examination of potential structural factors suggests that these explanations may be quite important, our state-level panel regressions suggest a large cyclical component, explaining perhaps as much as one third of the decline. The direct examination of workers who appear to have left the labor force out of discouragement supports a more modest, albeit still significant, role for the business cycle in explaining the drop in participation.

In section IV we present an alternative, more unified accounting for these (and other) factors, building on a model of the participation rate that had its genesis in the earlier Brookings paper mentioned above. This model attempts to simultaneously capture the contributions of aging, the business cycle, other measurable factors—such as changes in life expectancy, educational attainment, Social Security generosity, and marriage and fertility rates—and birth-cohort-specific factors that we have not so far identified. The model estimates that most of the recent decline in participation is structural, although the model's separation of the decline into cyclical, structural, and residual components depends somewhat on the time period of estimation and other details about the model's specification.

Combining the results from these different approaches, we find a range of possible estimates for the importance of cyclical and structural factors. In section V, we consider the implications of our analysis for the participation rate both over the next few years and over the next decade. Our analysis leaves some uncertainty surrounding how much the participation rate is currently depressed below its trend. Nevertheless, weighing the strengths of these various approaches, along with additional evidence that the pool of potential labor force re-entrants might not be very large, our overall assessment is that most—but not all—of the decline in the labor force participation rate since 2007 is structural in nature.

Consequently, in the near term, policymakers should not expect the participation rate to show a substantial increase from current levels. Looking farther out, over the next decade our model projects further declines in the

trend participation rate, primarily due to the continued transition of the baby boom generation into retirement.

I. Alternative Views of the Recent Decline in Labor Force Participation

The prominence of the decline in the labor force participation rate since 2007, along with its importance for policymakers, has fueled a substantial flow of recent research and commentary on this topic. While an extensive review of this literature is beyond the scope of this paper, as a prelude to our own analysis we provide a brief summary of some of the research focused on the recent decline in participation, highlighting the wide range of often contradictory conclusions reached in these studies.¹

In our view, observers should not have been particularly surprised by the fact that the labor force participation rate has declined noticeably over the past seven years. As noted above, our earlier Brookings paper, which was written prior to the financial crisis, had highlighted a number of factors likely to put downward pressure on labor force participation over the subsequent decade, and indeed, as shown by the dashed line in figure 1, the predictions we made in that paper turned out to track the decline in the actual participation rate well. That said, we readily admit that the severe recession complicates the interpretation of the participation rate decline, and, more generally, we would advise against taking an overly strong signal about the sources of the decline in the aggregate participation rate from our previous forecast. In particular, although the traditional view on movements in labor force participation over the business cycle has generally emphasized the absence of a substantial cyclical response, the breathtaking drop in labor demand in 2008 and 2009 may mean that this time really is different. In fact, the severity of the Great Recession and the subsequent slow pace of the economic recovery have led some researchers to interpret the decline in participation as having a large cyclical component.

A recent paper by Christopher Erceg and Andrew Levin (2013) provides a prominent example of this line of thought. Erceg and Levin first point out that labor force projections made by the Bureau of Labor Statistics (BLS) in November 2007 went badly off track over the next several years. This was true of both its projections for the aggregate participation rate and

1. For a more comprehensive survey on recent research on the decline in labor force participation, see Council of Economic Advisers (2014).

its projections for several key demographic groups.² They supplement that observation with a cross-state regression, showing a statistically significant negative correlation between changes in state-level participation rates for prime-age adults from 2007–12 and changes in state-level unemployment rates for this same demographic group between 2007 and 2010. Their conclusion from this analysis is that “the aggregate decline in prime-age LFPR can be *fully* explained by the persistent shortfall in labor demand” (p. 15), suggesting that the current level of the unemployment rate significantly understates the extent of labor market slack. While their analysis is suggestive, it warrants a number of caveats, including that their analysis covers a short time period and that it does not make use of information on the relationship between the unemployment and participation rates in previous episodes.³

At the other end of the spectrum, Marianna Kudlyak (2013) uses a bare-bones version of the cohort-based model we present later in the paper. She shows that the actual participation rate in 2012 was quite close to an estimate of the trend participation rate constructed from a model that includes only age-gender fixed effects and birth-year/gender fixed effects and that it was above the rate found by a model that takes into account the cyclical deviation of employment from its trend. Although she cautions that the estimated cohort effects may be influenced by both structural and cyclical factors, she interprets her results as suggesting that most of the decline in the participation rate is accounted for by the trend. However, it is difficult to assess Kudlyak’s interpretation, because

2. Using the BLS projections of the labor force participation rate from November 2007 as a baseline seems somewhat dubious to us, since the BLS projections of the trend through 2014 were well above those from our 2006 paper. The BLS projections for specific demographic groups are not projections from a behavioral model, but rather extrapolations of a nonlinear filter used to smooth historical labor force participation rates for each age, gender, race, and ethnicity group (see Toossi 2011). However, Toossi also reports on her preliminary efforts to construct a behavioral model for projecting the participation rate, by which she found that the projected values from such a model for the 2007–09 period were similar to those from the existing BLS model and that both approaches were surprised by the low level of the participation rate in 2009.

3. Additionally, as we demonstrated in an earlier version of this paper (Aaronson and others 2014a), Erceg and Levin’s conclusions rely crucially on the specification of their cross-state regression equations. We view the alternative panel data specifications that we use in this paper as more flexible. Another paper in this vein is by Hotchkiss and Rios-Avila (2013), who argue “that the dramatic decline in labor force participation during the Great Recession is *more than* explained by deteriorating labor market conditions (cyclical factors)” (p. 257). See Aaronson and others (2014a) for further comments on that work as well.

she does not include other factors that might cause changes over time in the propensity of different demographic groups to participate in the labor force, as we do below.

Other authors come out somewhere in the middle. For example, Daniel Aaronson and others (2014b) estimate a model that allows cohort effects and the coefficients on other controls to differ by age, sex, and educational attainment, and they find that more than half of the decline in the participation rate since 2007 reflected trend factors. Similarly, the Council of Economic Advisers (2014) attributes half of the decline to aging, one-sixth to “typical” cyclical weakness, and the remainder to other preexisting trends or other factors associated with the severity of the recession. A separate analysis by Robert Hall (2014) comes to a similar conclusion, although it traces much of the decline beyond that caused by aging to a combination of an increase in disability recipients and the expansion of the food stamp program, both of which discourage participation by implicitly taxing earnings. Finally, pure time-series methods, such as those employed by Willem Van Zandweghe (2012); Michelle Barnes, Fabià Gumbau-Brisa, and Giovanni Olivei (2013); and David Reifschneider, William Wascher, and David Wilcox (2013), attribute between half and two-thirds of the decline in participation since 2007 to trend movements, although of course such analyses say little about the underlying sources of a declining trend participation rate.

All these research papers provide a useful perspective on recent changes in the labor force participation rate. However, as Kudlyak (2013) concludes in her paper, “More research is needed that would explicitly model and account for the factors that influence the labor force participation decision of different demographic groups” (pp. 40–41). In particular, we think the most promising approach to analyzing participation rate movements would ideally incorporate insights from the voluminous literature on the factors that affect the labor force participation rate. This is what we attempt to accomplish in this paper.

II. Assessing the Importance of Structural Factors for the Decline in Participation

II.A. Aging and Retirement

The determinant of the aggregate participation rate that is perhaps the easiest to analyze is the changing age distribution of the population. As is well known, the population as a whole has been aging, putting downward pressure on the participation rate as the large baby-boom generation moves

Table 1. Estimated Contributions of Population Aging to Change in Aggregate Labor Force Participation Rate (LFPR) (Percentage Points)

	2007Q4 to 2014Q2	2000Q4 to 2007Q4	1990Q4 to 2000Q4	1976Q4 to 1990Q4
Actual LFPR	-2.8	-1.3	+0.2	+4.5
Contribution of aging	-1.3	-0.7	-0.1	+0.7

Source: Authors' calculations using data from the U.S. Census Bureau and the Bureau of Labor Statistics (Current Population Survey). Data are adjusted by the authors as described in the text.

into age groups that traditionally have low participation rates. This influence is particularly relevant for assessing the recent decline in participation, because the leading edge of the baby-boom generation reached age 62—the minimum age to receive Social Security retirement benefits—in 2008, coincident with the onset of the recession. Thus, we would have expected an upswing in retirements even absent the recession, which complicates efforts to distinguish structural and cyclical factors.

The first row of table 1 shows changes to the participation rate over different periods. These are calculated using Current Population Survey (CPS) microdata that have been seasonally adjusted and modified to account for periodic changes to population controls and the redesign of the CPS in 1994.⁴ (See section IV for more details.)

Row 2 of the table shows a “shift-share” calculation of the contribution of aging to the changes in the labor force participation rate (LFPR). In order to avoid any potential sensitivity to the choice of baseline for the calculation, we employ a chain-type calculation in which the age-specific participation rates are held constant only month by month.⁵ Specifically,

(1) Contribution of aging =

$$\sum_{age,sex} \sum_{m=Oct\ 2007}^{Jun\ 2014} \frac{LFPR_{age,sex,m+1} + LFPR_{age,sex,m}}{2} (share_{age,sex,m+1} - share_{age,sex,m}).$$

4. The change in the participation rate between 2007Q4 and 2014Q2 calculated using our adjusted data differs slightly from the published rate: the latter declined 3.1 percentage points between 2007Q4 and 2014Q2, compared to 2.8 percentage points in our data. We will refer to the adjusted-basis aggregate rate throughout the paper, although several of the analyses will use unadjusted detailed data.

5. Alternatively, we could have used a more familiar formula that holds either age-specific participation rates or population shares constant at some base year's values. The estimated contributions for the most recent period are similar using the more familiar formulas.

By this calculation, aging contributed 1.3 percentage points to the total decline of 2.8 percentage points in the aggregate participation rate.⁶ That is, nearly half of the observed decline can be attributed to the changing age distribution of the population.⁷ (Of course, as has been noted elsewhere, there may have been other demographic changes to the potential labor force offsetting some of this decline.⁸ We attempt to quantify the contributions of these and other factors in section IV.)

The importance of aging to the decline in participation is confirmed by the data on retirements shown in figure 2. The share of the working-age population reported as retired in the CPS (the solid line in the upper panel) has risen by more than a percentage point since 2007,⁹ just a bit below the contribution of aging reported in table 1.¹⁰ We can isolate the effects on retirement from the changing age composition of the population by holding the fractions retired at each age constant at their 2007 levels. This contribution, shown by the dashed line in the upper panel of figure 2, is similar to the results in table 1.

History suggests that the effect of the increase in retirements on labor force participation is likely to be persistent, because retirees in the past have tended to remain out of the labor force. As shown in the upper panels of figure 3, the probability that a CPS respondent observed to be retired in

6. This calculation actually represents the contributions of changes in the age-sex distribution of the population. However, the contributions of changes in the population's male-female composition are so small that we refer to our calculations simply as the contributions of aging.

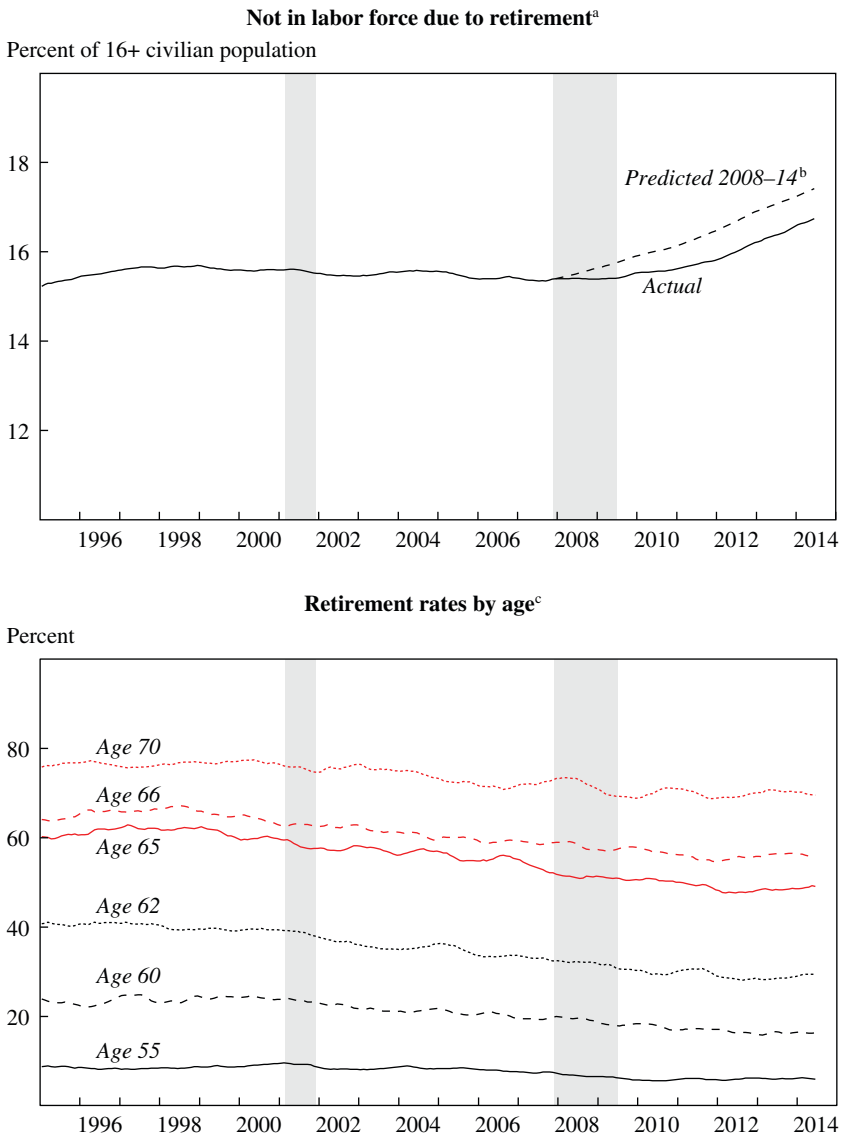
7. Note that the aging of the population has two major components. One is the movement of the large baby-boom cohorts from middle to more advanced ages. The other is the ongoing increase in longevity, which would tend to skew the age distribution toward older ages even if all birth cohorts were the same size. However, over this period the latter accounted for only a small portion of the contribution of the changing age distribution calculated here.

8. One could perform a similar exercise along several other dimensions, such as educational attainment or marital status, both of which would contribute positively to the change in the participation rate in recent years and thus offset some of the effects of population aging (see, for example, <https://sites.google.com/site/robertshimer/cbo-employment.pdf>).

9. The CPS-based estimates of retirement are also broadly in line with the data on Social Security retirement recipients, although the latter data show an upturn during the recession, perhaps because some individuals who lost their jobs began to collect Social Security retirement benefits during the recession, but nevertheless remained in the labor force.

10. Note that this exercise is a lower bound for the effects of aging, since retirement is only one component of age-related declines in participation (as participation rates are also lower for older workers who are not yet retired, for instance). Fujita (2014) also notes that CPS retirements have edged up over recent years, but he has not investigated whether this rise is consistent with aging and institutional changes as we do here.

Figure 2. Retirement Rates, Total and by Age Bracket, 1995–2014



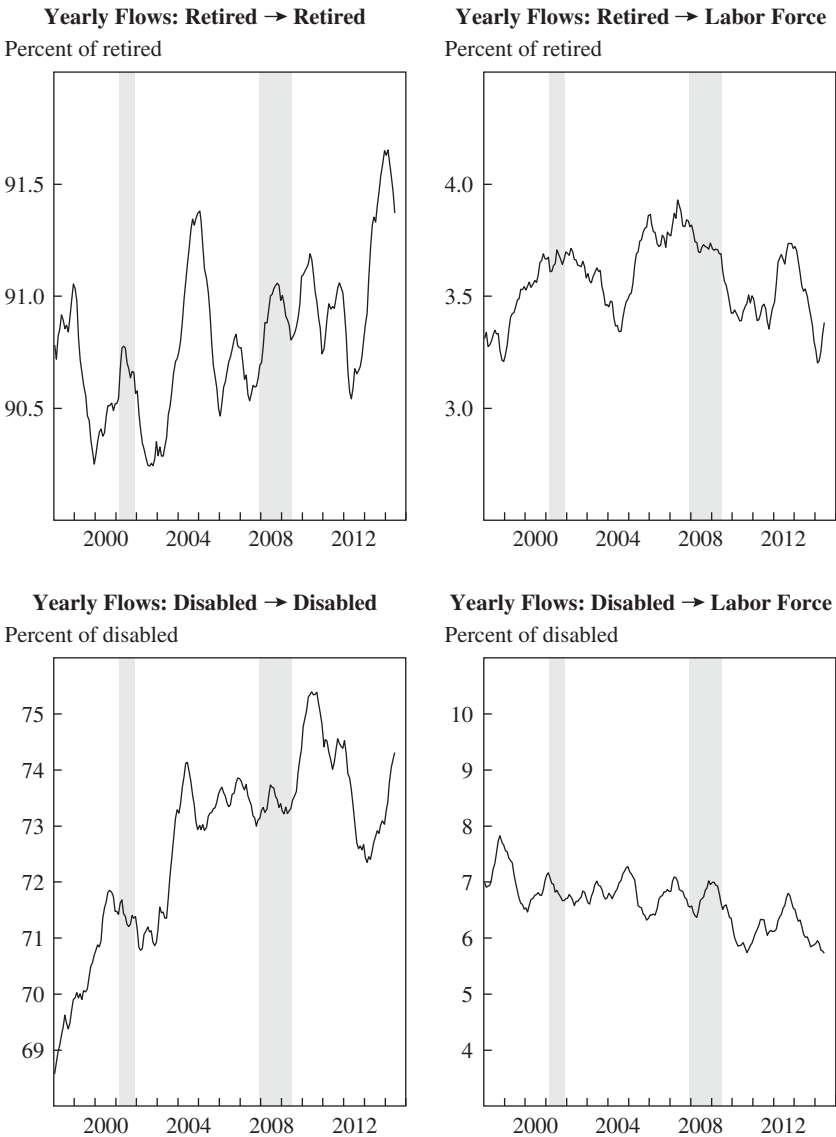
Source: Authors' calculations using data from the U.S. Census Bureau and the Bureau of Labor Statistics (Current Population Survey).

a. 12-month moving averages.

b. Within-age retirement rates fixed at 2007 levels. See text for more detail.

c. Percent of people of a given age who are retired; 12-month moving averages.

Figure 3. Retirement and Disability Flows 1998–2014^a



Source: Authors' calculations using data from the Current Population Survey Longitudinal Population Database (see Nekarda 2009).

a. 12-month moving averages.

one month is then also retired one year later is greater than 90 percent (left panel), and the probability that such a respondent is back in the labor force one year later is less than 5 percent.¹¹

That said, as measured by the CPS, actual retirements during the recession and early recovery have lagged what would have been predicted by age alone. Instead, within-age retirement rates, shown in the lower panel of figure 2, declined modestly on net, offsetting a portion of the effect of aging during this period. Should these offsetting declines be attributed to the recession? The decline in within-age retirement rates actually started around the late 1990s, likely due to a combination of institutional changes in social security and pension plans, increasing levels of education among older individuals, and longer life spans.¹² There is no obvious impact of the Great Recession and its aftermath on within-group retirement rates, and thus no clear evidence in favor of a dominant effect from either increased early retirements due to labor market discouragement in the wake of the recession or deferred retirements associated with the declines in household wealth incurred during the financial crisis.¹³

Overall, the CPS data are consistent with aging making an important contribution to the decline in the participation rate, and the fact that retirement has apparently not been distorted much by the Great Recession supports a structural interpretation of this phenomenon. Nevertheless, to the extent that individuals have shown any unusual cyclical movement out of the labor force and into retirement, it would have only a temporary effect on the trend. By age 66 nearly 60 percent of individuals are retired, and by age 70 about 70 percent are retired (lower panel of figure 2). Hence, as people age, what might have started as a premature retirement turns into an

11. These 12-month labor force flows as well as those presented later in the paper for the disabled and those who report themselves as wanting or not wanting a job relate to the flows between five states: in the labor force, not in the labor force but want a job, retired, disabled, and all other not in the labor force and who do not want a job. They are calculated by the authors from the Current Population Survey Longitudinal Population Database (see Nekarda 2009).

12. See Mastrobuoni (2009) and Blau and Goodstein (2010).

13. A recent study by Helman and others (2014) finds that the share of retirees saying that they retired earlier than planned rose from just under 40 percent in the years prior to the recession to closer to 50 percent during the recession and in the years since. The share reporting retiring later than expected also edged up slightly. Gorodnichenko, Song, and Stolyarov (2013) find that over time white men have increasingly reacted to recessions by retiring. While this evidence is suggestive of a behavioral response to the recent business cycle, it is difficult to see in the aggregate data from the CPS. Bosworth and Burtless (2010) also find economically modest effects of the Great Recession on the labor force participation rate of older workers.

expected retirement, and the participation rate converges to its trend. This is not to say that early retirement has no cost to society. But its implications for the aggregate labor force participation rate over time are muted.¹⁴

II.B. Teenagers and Younger Adults

As highlighted in the top panel of figure 4, the labor force participation rate for 16- to 24-year-olds declined from about 59 percent in 2007 to roughly 54½ percent in mid-2014. In fact, the decline for this group has been so large that it accounts for roughly 20 percent of the decline in aggregate participation over this period.¹⁵ However, as has been documented by a number of researchers, the decline in labor force participation in these ages long predates the Great Recession. Therefore, although a cohesive explanation for this downward trend remains elusive, we consider it structural.¹⁶ Nonetheless, the magnitude of the drop since the onset of the recession makes it worthwhile to explore whether there has been an unusual cyclical component to participation rates for this age group in recent years.

We begin with an examination of the extent to which an increase in schooling or schooling intensity could explain the decline in participation for this age group. Consistent with education as an important explanation, figure 5 shows that for both high-school (upper panel) and college-aged (lower panel) persons, the drop in labor force participation among youths over the past few decades coincides with a general rise in their school enrollment rates; likewise, youths who report being enrolled in school in the CPS have lower labor force participation rates, on average, than non-enrollees.¹⁷ However, the figure also shows that participation rates have

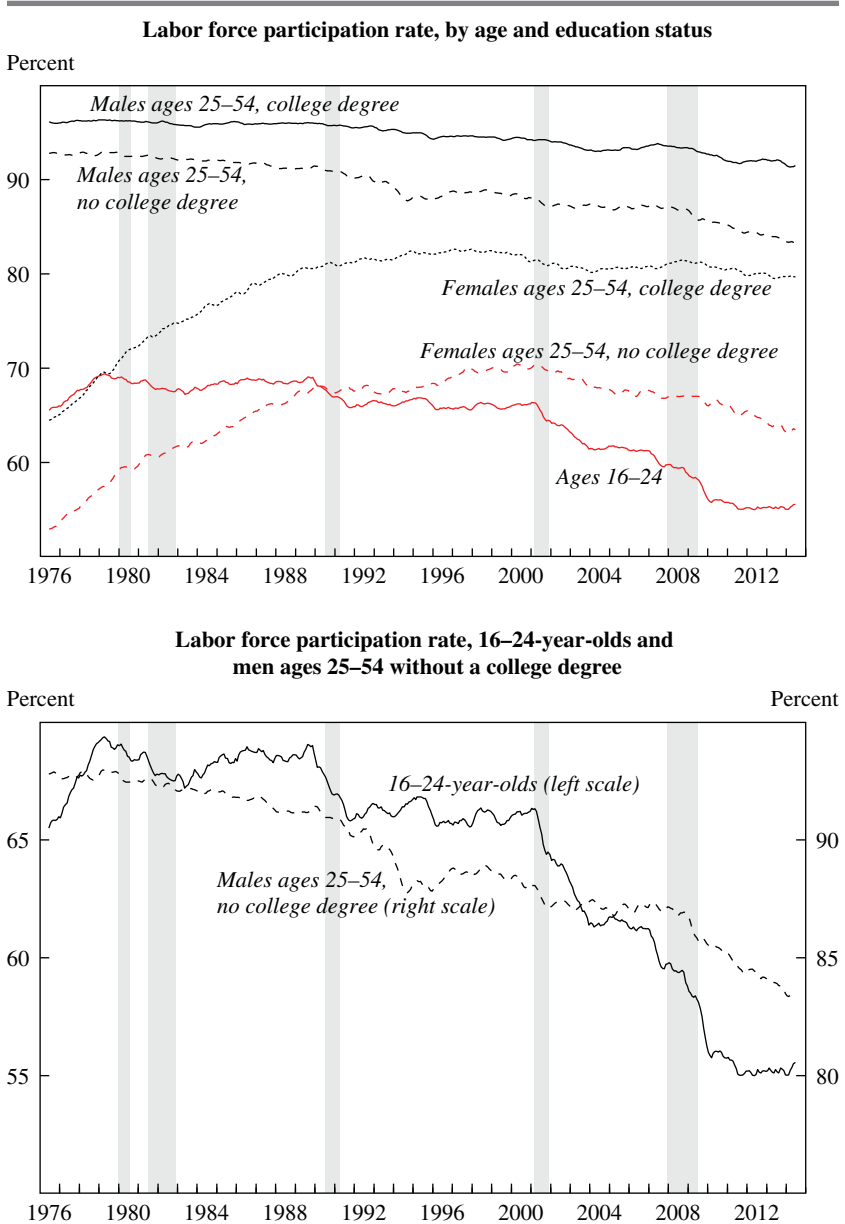
14. The same is also true, to a lesser extent, for exits into disability, since these are more prevalent among people close to retirement age.

15. For comparison, 16- to 24-year-olds accounted for less than 15 percent of the labor force in 2007.

16. For some earlier work on the subject, see Aaronson, Park, and Sullivan (2006); Morisi (2008, 2010), and Smith (2011).

17. In these figures, enrollment is measured by the response to the CPS question: "Last week was [the respondent] attending or enrolled in a high school, college, or university? Yes if currently on holiday or seasonal vacation, no if on summer vacation." Of note, the question is phrased such that persons on summer vacation should not report themselves as being enrolled. Unfortunately, it appears that the enrollment variable in the monthly CPS after 2012 is not strictly comparable to that in prior years, which we suspect may be due to changes in the way the question was asked. Prior to 2013, current enrollment status was only asked for respondents up to age 24, but starting in 2013 it was asked of all respondents under 54 years old. Coincident with this change, reported enrollment for ages 24 and younger saw a discrete drop in January 2013 and thereafter. For this reason, we only show enrollment rates through 2012, and our counterfactual activity only examines changes in participation through 2012.

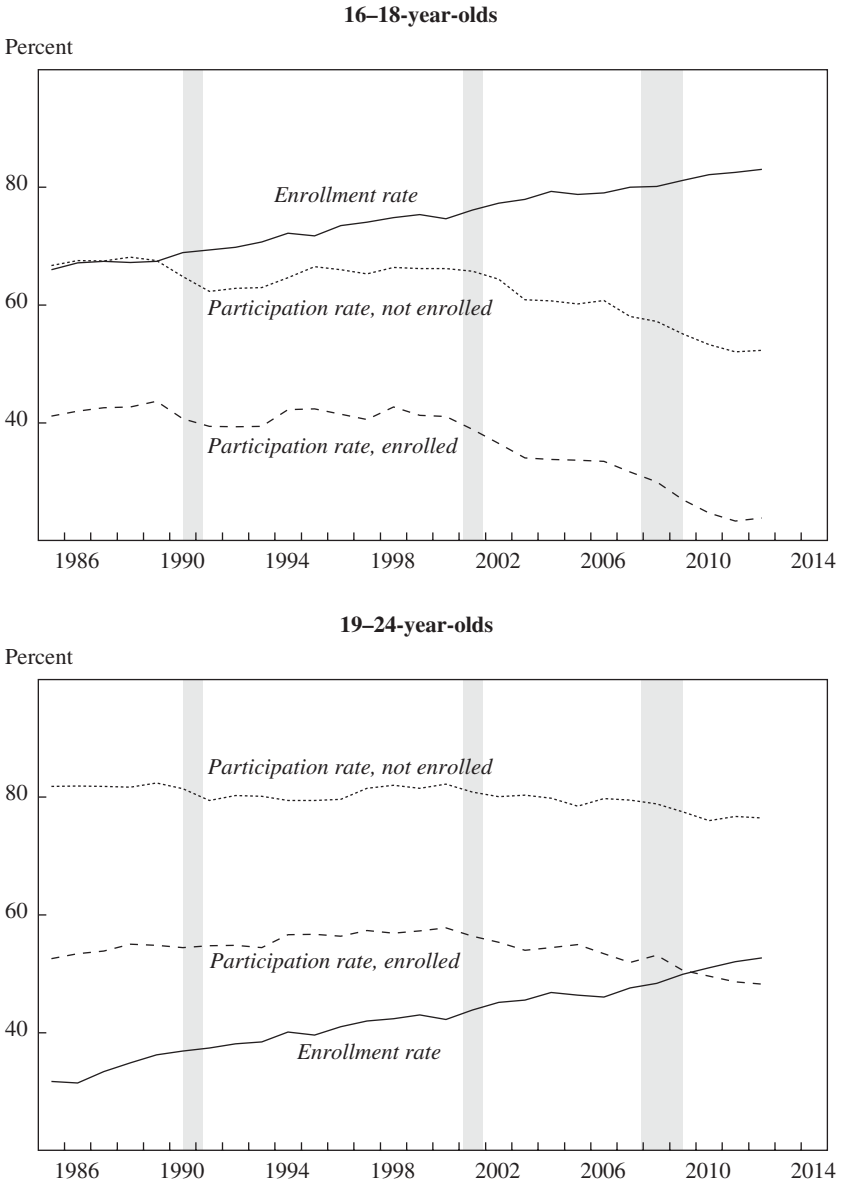
Figure 4. Select Participation Rates by Age, Gender, and Education 1976–2014^a



Source: Authors' calculations using data from the U.S. Census Bureau and the Bureau of Labor Statistics (Current Population Survey).

a. 6-month moving averages; data are seasonally adjusted by the authors.

Figure 5. Labor Force Participation and School Enrollment, 16–24-Year-Olds 1985–2012^a



Source: Authors' calculations using data from the U.S. Census Bureau and the Bureau of Labor Statistics (Current Population Survey).

a. Annual averages.

been falling for *both* enrolled and non-enrolled individuals, indicating that rising enrollment cannot explain the entirety of the overall decline.

The evolution of participation among young people is likely heavily influenced by an increase in the college earnings premium and the return to skill more generally. Among enrollees, a portion of the decline may be attributable to increased educational intensity—such as greater time spent on homework or other extracurricular activities—since there is a greater return to such activities.¹⁸ The decline in participation among the non-enrolled could also be linked to this trend. One effect of the rising return to education and to skill more generally has apparently been a decrease in demand for some adult workers lower on the educational scale. This drop in demand may have displaced them into lower-skilled sectors, thus increasing competition for jobs in the low-skilled labor market (such as retail sales and food service) and crowding out younger job seekers. Another ongoing source of crowd-out may be the increasing population share of less-educated adult immigrants, with some evidence suggesting that the displacement effect of immigration is much larger on the employment of younger persons than on the employment of prime-age adults (Smith 2012).

Figure 5 shows that recent increases in enrollment are mostly a continuation of these past trends, which, as we have argued, appear to have largely structural explanations. A shift-share exercise suggests that the rise in enrollment can explain about one-fourth of the decline in participation for 16- to 24-year-olds between 2007 and 2012, or about 0.2 percentage point of the decline in the aggregate participation rate. Likewise, participation rates among both enrollees and non-enrollees have generally fallen in line with their trends since the late 1990s; together they account for about 0.6 percentage point of the decline in the aggregate participation rate between 2007 and 2012.¹⁹ However, while the low rates of youth

18. For high-school-age persons, this is consistent with the findings in Ramey and Ramey (2010), which show that the amount of time spent by parents (especially college-educated parents) on activities for their children has been rising over time (particularly for older children and in the “travel” and “activities” categories).

19. To arrive at this estimate, we decomposed the decline in participation from 2007 to 2012 into: (i) the decline attributable to falling participation for enrolled persons (the change in participation for enrolled persons multiplied by the share enrolled in 2007); (ii) the decline attributable to falling participation for non-enrolled persons (the change in participation for non-enrolled persons multiplied by the share not enrolled in 2007); and (iii) the decline attributable to the rise in the enrollment share (difference in participation for enrolled and non-enrolled persons in 2012 multiplied by the average change in enrollment rates from 2007 to 2012).

participation seem to us to be primarily a function of long-run trends, the associated higher levels of education should raise the participation rates of today's young people as they enter their prime years and beyond, perhaps providing some boost to aggregate participation as well.

II.C. Less-Educated Adults

Returning to figure 4, the recent decline in participation among prime-age males without a college degree also appears to be the continuation of a long-term secular decline. Moreover, since the early 2000s, men have been joined in this downtrend by prime-age women without a college degree. These declines have been the subject of a considerable body of literature reaching back to the 1980s. The early literature (for example, Juhn 1992), which focused on prime-age men, identified declining labor market opportunities for low-skilled workers, manifested in stagnant real wage growth, as the likely explanation.

Since the 1990s, changes in labor demand have apparently not been characterized by a monotonic decline in the demand for low-skilled workers but rather by a relative decline in labor demand for occupations that have tended to be “middle-paying,” or middle-skill jobs (Autor 2010), which we call below, “middle-type” jobs (referring jointly to the wage and skill dimension of the job).²⁰ Just as the decline in demand for low-skilled workers appears to have accounted for much of the decline in participation among men in the 1970s and 1980s, polarization in labor demand seems at least plausible as an explanation for some of the subsequent secular decline in participation among less-educated individuals. The idea is that polarization displaces some workers who had been employed in middle-type jobs. Of these, some are able to move to high-type jobs, some move into the lower-paying service sector, and some may drop out of the labor force. In addition, the movement of workers from middle-type jobs into lower-paying occupations may also displace lower-skilled workers, leading some of them to drop out of the labor force as well. Labor force withdrawal is likely to be most acute for less-educated

20. The latest research explains labor market polarization as a consequence of two factors: the decline in the cost of computing and automation technology and the increased accessibility of overseas labor and product markets (see Autor, Dorn, and Hanson 2013). While these developments appear to have directly reduced labor demand for individuals in middle-type occupations, they likely *raised* labor demand for more educated individuals and had little direct effect on labor demand for service-sector jobs.

adults because they are most likely to have been employed in middle-type or lower-type jobs.²¹

Consistent with the hypothesis that polarization can account for the secular decline in participation, we find that since 1985 participation rates for less-educated adults fell further in states with a greater decline in middle-type employment shares.²² Additionally, although polarization is a long-term phenomenon, some recent research suggests that it accelerates during recessions (Jaimovich and Siu 2012). Indeed, the participation rate for less-educated prime-age men has exhibited a stair-step pattern, discretely falling during recessions and failing to recover thereafter (figure 4, lower panel). Whether this pattern represents an actual acceleration of the displacement process or just the re-emergence of the trend during a cyclical downturn, the result is the same—these men do not reenter the labor force as the economy expands. Thus, to the extent that polarization explains some of the decline in labor force participation both before and during the Great Recession, we would not expect to see an increase in participation among these workers as the unemployment rate moves back toward its natural rate.²³

II.D. Disability

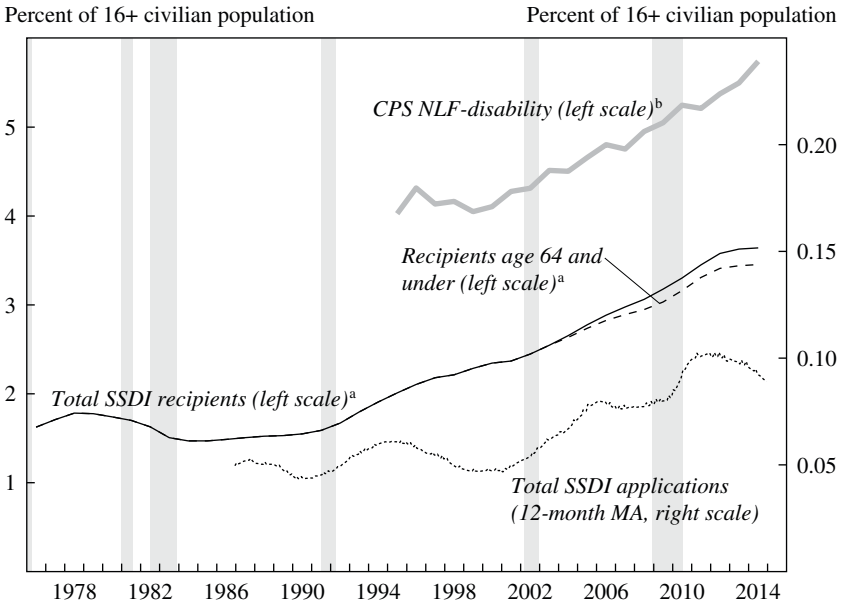
Another possible explanation for the decline in participation rates among adult men and women is the expansion of disability insurance. The gray solid line in figure 6 depicts CPS self-reported nonparticipation in the labor force due to disability. As a percent of the U.S. working-age population, nonparticipation due to disability has been edging up by about 0.1 percentage point annually over the past decade or so, with only a little evidence of cyclicity.²⁴ Moreover, nonparticipants who report being

21. In 1985, of those without a four-year college degree, 21 percent of prime-age males and 45 percent of prime-age females were employed in middle-type jobs. Of those with a four-year college degree or more, only 11 percent of males and 28 percent of females were employed in middle-type jobs.

22. See Aaronson and others (2014a). This is also consistent with evidence from Foote and Ryan (2014), which used variation in participation and employment shares by job type across age, education, and Census divisions to show that there is a strong negative relationship between participation and the probability (if employed) of being employed in middle-type jobs.

23. This discussion is not meant to preclude other possible explanations for the decline in participation among prime-age workers, prime-age men in particular. One alternative is the increase in labor market opportunities for females, which could have resulted in a shift in the relative gender balance for providing household income.

24. For a further analysis of factors behind the rising disability, see Autor and Duggan (2003); Duggan and Imberman (2009); and Autor (2011).

Figure 6. Rates of Disability Recipiency and Applications, 1976–2014

Source: Authors' calculations using data from the U.S. Census Bureau and the Bureau of Labor Statistics (Current Population Survey) and from the Social Security Administration.

a. SSDI recipients data are end-of-year.

b. CPS not-in-the-labor-force-disability data are for December.

disabled are generally unlikely to return to the labor force. As shown in the lower panels of figure 3, the probability that a nonparticipant who is disabled will also report being disabled a year later is greater than 70 percent, and the probability that such a person will be in the labor force a year later is only 6 to 7 percent.

While the CPS definition of disability does not depend on the receipt of Social Security Disability Insurance (SSDI), and there is a difference in the levels of disability as measured by the CPS data and the administrative data on the receipt of SSDI benefits (the black solid line in figure 6), their trajectories are similar.²⁵ Focusing on the administrative data, once we limit our purview to persons ages 64 and under in order to abstract from the increase in SSDI receipts resulting from the rise in the Social Security

25. As pointed out in Fujita (2014, p. 3), although CPS disability is self-reported and is not related to the receipt of SSDI, it is nonetheless a fairly strict definition.

full retirement age from 65 to 66, there is only scant indication of a cyclical increase in benefit receipt—too small, as a share of the population, to have had much impact on the aggregate participation rate.

Applications for SSDI (dotted line in figure 6) nevertheless did step up during the recession and only began to ease in late 2012. This increase in applications alone could have resulted in a decrease in labor force participation, since applicants often view themselves as effectively disqualified from working (Autor 2011). However, we would expect any such effect to be reflected in the CPS disability measure, which, as mentioned, does not seem to have risen more quickly after 2008 than before. To the contrary, disability in the CPS seems to have continued to rise in recent years, even as applications and actual benefit receipts have stabilized, a discrepancy for which we have no ready explanation.²⁶

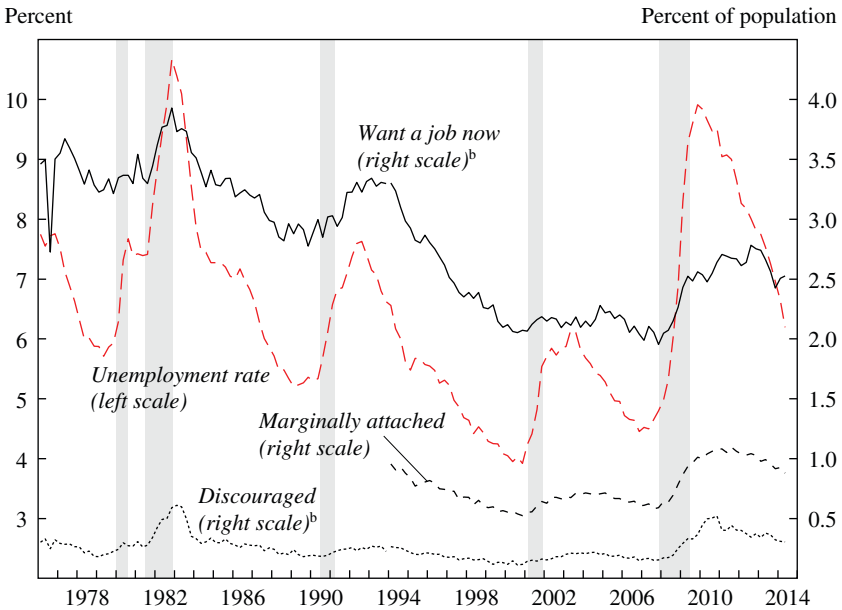
III. Assessing the Recent Cyclicity of Labor Force Participation

The bulk of the evidence presented so far suggests that structural factors can explain most of the decline in participation since 2007. However, we must also consider the possibility that the deep recession and slow pace of recovery led an unusually large number of persons to temporarily drop out of the labor force in recent years because they were discouraged about their job prospects. If so, these persons could return to the labor force when economic conditions improve sufficiently. In this section we assess how much the participation rate appears to be cyclically depressed, using a number of different approaches.

III.A. Reasons for Nonparticipation

We begin with an examination of the reasons that labor market nonparticipants give for being out of the labor force when responding to the CPS, which asks nonparticipants several questions aimed at identifying discouragement. First, a nonparticipant reports whether he or she wants a job. If, in addition, he or she is currently available to work and has looked for a job in

26. The modest cyclicity of nonwork resulting from disability is likely due in part to the fact that the criteria for being disabled do not change with the cycle. So while people are more likely to apply for disability when the economy is weak, actual reciprocity increases less. On the other hand, when the labor market is particularly strong, fewer people apply for disability benefits, opting to remain in the labor force, although the number of people involved is small relative to the size of the civilian population over age 16.

Figure 7. Rates of Selected Not-in-Labor-Force Categories and the Unemployment Rate, 1976–2014^a

Source: Authors' calculations using data from the U.S. Census Bureau and the Bureau of Labor Statistics (Current Population Survey).

a. Seasonally adjusted by authors (except post-1994 "Want job" and the unemployment rate series, which are seasonally adjusted by the Bureau of Labor Statistics).

b. Pre-1994 "Want job" and "Discouraged" series are based on outgoing CPS rotation groups only.

the past 12 months, the BLS classifies him or her as "marginally attached." If he or she also has a job-related reason for not currently looking for work, then he or she is classified as "discouraged."²⁷ While these individuals are not officially counted as being in the labor force, the data on gross labor market flows indicate that they have significant probabilities of moving into the labor force (we discuss the evidence on gross labor market flows data in section V).²⁸

As indicated in figure 7, these three successively stringent measures of labor market discouragement did increase during the Great Recession and

27. While only persons who satisfy all four criteria are classified by the BLS as "discouraged workers," we will take a more expansive view here.

28. The discussion here and in the remainder of this section uses data that have not been adjusted for changes to population controls.

the early stage of recovery, consistent with temporary labor market withdrawal by individuals who faced poor job prospects during that period. However, the share of the population in each of these categories is relatively small, ranging from 0.3 percent for discouraged workers to slightly below 1 percent for the marginally attached and to 2½ percent for those who say that they want a job (which includes the discouraged and marginally attached). Indeed, judging by the behavior of the broadest (“want a job”) category, the rise in discouragement at its peak in 2012 could explain at most ¾ percentage point of the decline in the labor force participation rate since the end of 2007.²⁹

The “want a job” category appears to have lagged somewhat behind the unemployment rate during the Great Recession, peaking three years later at the end of 2012. This lagging behavior can also be observed in prior recessions and suggests that the cyclical component of labor force participation might lag changes in the unemployment rate, an issue that we investigate more formally below. More recently, all three measures of discouragement have declined somewhat, although they remain elevated, suggesting that the labor force participation rate was cyclically depressed in 2014Q2 by perhaps ½ percentage point due to labor market discouragement as most broadly defined.³⁰

Of course, some amount of discouragement may manifest itself in ways that are unlikely to show up in these measures. For example, business-cycle conditions almost certainly affect individuals’ decisions to enroll in school, apply for disability insurance, retire, or stay home and take care of house

29. This is based on a counterfactual of no change in the share of civilian working-age population reporting themselves wanting a job. This counterfactual seems reasonable based on the pattern earlier in the decade, although over the longer term the share has been trending downward (Barnichon and Figura 2013).

30. Simple regressions of these measures from 1994 to 2007 on the unemployment rate gap and its lag suggest that the fractions of marginally attached and discouraged workers are slightly above the levels one would expect given the level of the unemployment rate. We would caution, however, that measurement error may be more acute for these measures of labor market discouragement than with the more familiar measures of unemployment and labor force participation, because these more-detailed not-in-labor-force classifications depend on subjective criteria. In particular, labor market conditions may affect how respondents answer survey questions that are used to determine their status within the nonparticipation category. For example, if wage rates were rising more quickly, more nonparticipants would likely claim that they want a job, indicating that the measures in figure 7 might be understating the true extent of labor market discouragement. In addition, Barnichon and Figura (2013) argue that the share of want-a-job nonparticipants might have an important secular component, thus making the inference of cyclicalities in labor force participation from this category of nonparticipants even more difficult.

or family. And many of these individuals may not report themselves as wanting a job. As we discussed previously, with respect to disability, retirement, and schooling, much of the recent decline appears to be structural; nevertheless, determining how much of these movements reflects cyclical factors and how much reflects structural factors is difficult. We will return to some of these issues later in the paper.

III.B. State-Level Panel Data Analysis

An alternative way to assess the cyclicity in labor force participation is by exploiting both time and geographical variations in labor market conditions and the participation rate. In particular, the timing and severity of business-cycle fluctuations typically vary considerably across states, which provides an additional source of identification that can potentially help to estimate dynamic linkages between unemployment and participation. Of course, this identification also has its limitations in that adequate state-level data are not available to explicitly control for many of the structural influences we highlighted above. We attempt to capture these factors through the inclusion of state-specific linear time trends, although they may be incomplete controls.

We estimate state-level panel regressions on annual data using the following specification:

$$(2) \quad LFPR_{s,t} = \alpha_s + \delta_t + \gamma_s time + \lambda LFPR_{s,t-1} + \sum_{i=0}^k \beta_i UR_{s,t-i} + \phi \mathbf{X}_{s,t} + \varepsilon_{s,t},$$

where α_s represents state fixed effects, δ_t are time fixed effects, γ_s are state-specific time trends, and \mathbf{X}_s is a vector of covariates related to demographics.³¹

As shown in column 1 of table 2, the cyclical parameter β_0 is -0.16 when only the contemporaneous unemployment rate is included in the

31. For these specifications, we construct labor force participation rates, unemployment rates, and demographic controls by state-year from the CPS microdata. The covariates included in vector \mathbf{X}_s are the share, by state and year, in each of 24 demographic groups defined by sex, education, and age (where the two education groups are persons with no more than a high school degree and persons with at least some college or more, and the six age groups are 16–24, 25–34, 35–44, 45–54, 55–64, and 65 and older). Estimates from specifications that excluded state-specific trends were qualitatively similar to the results shown in table 2. We also experimented with state-level labor force participation rates and unemployment rates from the BLS's Local Area Unemployment Statistics (LAUS); since the latter data are very close to the CPS microdata estimates at the *annual* frequency, the obtained results were also qualitatively similar to the ones reported here.

Table 2. Results from State Panel Regressions^a

<i>Dependent variable:</i>	(1)	(2)	(3)	(4)	(5)
<i>Labor force participation rate</i>	1978–2013	1978–2013	1978–2013	1990–2013	1978–2007
Labor force participation rate _{t-1}			0.38*** (0.03)	0.27*** (0.03)	0.34*** (0.03)
Cyclical parameters:					
Unemployment rate _t	-0.16*** (0.03)	-0.02 (0.04)	-0.04 (0.03)	0.01 (0.04)	-0.06 (0.04)
Unemployment rate _{t-1}		-0.09*** (0.03)	-0.09** (0.03)	-0.08 (0.05)	-0.10*** (0.03)
Unemployment rate _{t-2}		-0.07** (0.03)	-0.04 (0.03)	-0.03 (0.04)	-0.04 (0.04)
Unemployment rate _{t-3}		-0.14*** (0.03)	-0.06** (0.03)	-0.10*** (0.04)	-0.08** (0.04)
Total cyclical effect	-0.16*** (0.03)	-0.32*** (0.03)	-0.37*** (0.04)	-0.27*** (0.07)	-0.43*** (0.05)
Implied cyclical shortfall in 2014Q2 (percentage points)	-0.11	-0.94	-1.03	-0.89	-1.14
No. of observations	1,836	1,683	1,683	1,224	1,377
No. of states	51	51	51	51	51

Source: Authors' estimates using data from the U.S. Census Bureau and the Bureau of Labor Statistics (Current Population Survey) and from the Congressional Budget Office (natural rate of unemployment).

a. All regressions are estimated using annual data and are weighted by state population. All specifications include state and year fixed effects, state-specific linear time trends, and demographic controls. Statistical significance indicated at the *** 1 percent and ** 5 percent level. Standard errors clustered at the state level are in parentheses.

specification, suggesting a rather modest cyclicity of labor force participation.³² Defining the cyclical shortfall in the participation rate in 2014Q2 to be β_0 multiplied by the cyclical shortfall in the unemployment rate as estimated by the CBO (the difference between the actual unemployment rate and the CBO's estimate of the long-run natural rate, a gap that was about 0.7 percentage point in 2014Q2), the estimated cyclical shortfall in the participation rate from this specification as of 2014Q2 was only -0.1 percentage point.

The next column adds three lags of the unemployment rate in order to test the hypothesis that the participation rate responds to labor market shocks with some delay.³³ These lags are all negative in sign and tend to be statistically significant at conventional levels.³⁴ The high autocorrelation of the unemployment rate gives rise to the classical problem of multicollinearity, and we would thus caution against interpreting each separate lag estimate with high confidence. Indeed, statistical tests show that lag estimates are statistically indistinguishable from each other. Nevertheless, taking the estimates at face value, the implied cyclical shortfall in participation in 2014Q2 is nearly 1 percentage point, substantially above the specification with only contemporaneous relationship.³⁵ This result suggests both that the labor market history matters for movements in participation rate and that prolonged recessions (or expansions) can have a stronger cyclical effect on participation than short-lived business-cycle swings.

We next add one lag of the participation rate to the regression specification. Theoretically, lags of the participation rate could matter if labor force participation decisions are persistent; for example, if hysteresis effects or

32. One caveat to the use of the unemployment rate as the indicator of the business cycle is that it might be endogenous to changes in the participation rate. An alternative indicator of the business-cycle available at the state level is the (detrended) level of employment. However, since (state) employment trends necessarily depend on (state) trend participation rate movements, the employment gap measures suffer from the same problem of endogeneity.

33. The Council of Economic Advisers (2014) reports estimates from similarly specified regressions using national data and finds significant lagged effects from the unemployment gap on participation of up to 8 quarters. In addition, the IMF (2014) and Aaronson and others (2014b) report estimates from similar regressions using state-level data and find lagged effects of the business cycle on participation up to 3 years.

34. When we included in our regression specification lags beyond 3 years, they tended to be statistically insignificant and quantitatively rather small.

35. The cyclical shortfall in 2014Q2 is calculated as described for the regression of column 1, except that we also account for lags by multiplying the coefficient on each lag by the second-quarter estimate of the cyclical shortfall in unemployment (actual unemployment rate less the CBO's estimate of the natural rate) for 1, 2, or 3 years previously.

transition costs are important.³⁶ If this were the case, our finding that lags of the unemployment rate are sizable could reflect the indirect effect of the cycle on the participation rate in previous years, rather than direct effects of lagged unemployment rates. Indeed, when we include one lag of the participation rate (column 3 in table 2), the estimated coefficients on the lags of the unemployment rate are somewhat smaller. Nevertheless, in this specification, the implied current cyclical shortfall in participation remains about the same as the previous specification (1 percentage point).³⁷

The final two columns of table 2 explore the robustness of these findings to different sample periods. In column 4 we limit the sample to 1990 and later, while in column 5 we limit the sample to 2007 and earlier. In both cases, the estimated cyclical shortfall is similar to the one obtained from estimates on the entire sample.

To summarize, our investigation of the relationship between state-level unemployment and participation rates over a multi-decade period suggests that the cyclical shortfall in the participation rate in 2014Q2 is between -0.1 percentage point (no lags of the participation or unemployment rates) and -1 percentage point (one lag of the participation rate and three lags of the unemployment rate), with most of the estimates at the upper end of the range. The associated estimates of the contribution of the cycle to the change in the participation rate since 2007Q4 are essentially the same, since the unemployment rate was very close to the CBO natural rate just prior to the recession.

The state-level regressions are appealing, because they provide additional variation for a business cycle phenomenon for which we typically have few observations, but the identification is not without its limitations. For instance, there may be spurious correlation between the unemployment rate and the participation rate due to measurement error. Even abstracting from measurement error, it is conceivable that, as mentioned previously, the correlation might be spurious due to an omitted variable. For example, states with a larger share of older population (and thus more subject to

36. See, for example, Clark and Summers (1982).

37. To estimate the cyclical shortfall implied by these regressions, we start by estimating the shortfall implied directly by the unemployment rate and its lags, as described previously. However, lags of the unemployment rate also contribute indirectly through effects on lagged participation. Since participation one year ago is also influenced by participation in the previous year (and hence, another three lags of the unemployment rate), the total cyclical contribution of the unemployment rate and its lags is an infinite series that is a function of the coefficients on the lagged participation and unemployment rates. In our calculation, we stopped the approximation at the third lag of the actual labor force participation rate.

participation declines due to aging), like Florida and Arizona, might be more prone to housing booms and busts and thus to more severe downturns. In addition, the magnitude of the cyclical effect is sensitive to the inclusion of lags of the unemployment rate, which gives us pause. The econometric evidence supports the inclusion of three lags of the annual unemployment rate, and it is true that in a recession individuals may make decisions about their labor supply that are sticky.³⁸ However, as noted previously, the literature has generally found that the participation rate on net is only modestly cyclically sensitive, and we provide further support for this finding later.³⁹ Because of these limitations, empirical models that explicitly control for some of the above-mentioned factors, such as the one we present in the next section, are better equipped to distinguish between trend, as we have defined it, and cycle.

We conclude this section by cautioning that given the depth of the recession and the slowness of the recovery in this episode, the particularly large cyclical shortfall in participation implied by these regressions may lead to several different outcomes. First, the participation rate changes could be related to discouragement that will reverse as the economy strengthens further. Second, the recession might have merely accelerated trend declines that would have happened anyway (for example, due to retirements) and hence are unlikely to reverse. Third, the recession may have caused cyclical declines in the participation rate that will eventually become permanent (for example, due to retirement or disability).

IV. Declining Participation through the Lens of a Model: Updating the Model from the 2006 Brookings Paper

While the analyses above go some way toward quantifying the extent to which aging, the business cycle, and an assortment of other factors can explain recent declines in labor force participation, their factor-by-factor approach is limited in its ability to decompose the aggregate decline into

38. For instance, decisions to attend school or care for children may entail multiyear commitments. Moreover, as noted by Autor (2011), the disability application process frequently lasts up to 3 years. When we ran state-level panel regressions for separate demographic groups, we found particularly large lagged effects of the unemployment rate in the regressions for low-skilled (at most high school) and older (55 years and over) individuals, consistent with disability as an explanation and with that of accelerated retirement.

39. One study that has tried to address this issue using individual panel data for white male college graduates, Kahn (2010) finds that entering the labor market in a poor economy has lasting effects on wages and occupational attainment, but not on labor supply.

cyclical and structural components in an integrated and consistent fashion. Also, outside of the contributions of aging and the cycle, it is difficult to use these analyses to project the path of aggregate labor force participation over the next decade. In this section, we turn to an alternative and complementary approach based on an updated version of the model we introduced in Aaronson and others (2006).⁴⁰

IV.A. Description of the Model

In this cohort-based, demographically disaggregated model, we combine the changing age distribution with various factors to explain within-age changes in participation. We refer the reader to the 2006 Brookings paper for a general description of the model and its motivation.⁴¹ However, we lay out the model briefly here, since the specification has changed significantly since 2006.⁴²

Form and estimation. The model has this form:

$$(3) \quad \log \left(\frac{lfp_{a,t,s}}{1 - lfp_{a,t,s}} \right) = A_{a,s} + K_{t-a,s} + \mathbf{X}_{a,t,s} \boldsymbol{\lambda}_{a,s} + \varepsilon_{a,t,s},$$

where lfp represents the seasonally adjusted labor force participation rate expressed as a fraction; a is age (in single years), between ages 16 and 79⁴³; t is calendar time (in quarters); s is sex; A is an age-and-sex-specific constant, that is, an “age effect”; K is a birth-year-and-sex-specific constant, that is, a “cohort effect”⁴⁴; \mathbf{X} is a vector of variables that may vary by age, time, and/or sex; $\boldsymbol{\lambda}$ is a vector of coefficients, which generally vary by both age and sex (some coefficients are constrained to be zero for some age-sex groups); and ε is an i.i.d. error term.

40. Stephanie Aaronson, Bruce Fallick, Andrew Figura, Jonathan Pingle, and William L. Wascher, “The Recent Decline in Labor Force Participation and Its Implications for Potential Labor Supply,” *Brookings Papers on Economic Activity*, Spring 2006.

41. In addition to the earlier Brookings paper, see Fallick and Pingle (2007).

42. Several of the changes adopted improvements introduced by Balleer, Gomez-Salvador, and Turunen (2009). See also Balleer, Gomez-Salvador, and Turunen (2014); Benito and Bunn (2011); Kawata and Naganuma (2010); and Duval-Hernández and Romano (2009).

43. We do not model the participation rates of persons 80 years or older because of the small sample size and very low participation rates of this group. In summing to an aggregate participation rate, we treat the rate of the 80+ group as always at trend.

44. Because birth dates are spread throughout the calendar year, current year minus reported age is not a perfect measure of birth year, nor should the cohort effects jump sharply from one birth year to the next. Therefore in the estimation each cohort effect K contributes in a weighted fashion to the equations for adjacent years.

We draw the data on participation rates by age and sex from the micro CPS files, adjust the raw series to account for changes in the survey and changes in the population controls, and seasonally adjust them.⁴⁵

We include 10 variables in the vector \mathbf{X} , as described next.

(1) The aggregate *unemployment rate gap*, divided into positive and negative components in order to allow for asymmetric responses to tight and loose labor markets. We use the long-term natural rate estimated by the CBO to define the unemployment rate gap. Our baseline specification includes the contemporaneous gaps and lags at 4, 8, and 12 quarters. We have also estimated the model using only the contemporaneous unemployment rate gap and with other lag lengths; all yield similar estimates of the contributions of trend and cycle to the decline in participation since 2007.⁴⁶

(2) The aggregate *personal bankruptcy rate*, as a percent of the population. In principle, household wealth should influence participation decisions, most notably for retirement. However, these effects are difficult to identify in aggregate data, possibly because holdings of wealth are so skewed. The personal bankruptcy rate is intended to represent changes in household wealth at the most relevant parts of its distribution.

(3) The percent of each age-sex group with a *college degree*. Participation rates differ significantly by education, which is typically attributed to a combination of higher returns to market work and, especially at older ages, the lower physical demands usually associated with occupations requiring greater education. Although we include only college attainment explicitly, this variable is intended to represent the patterns in educational attainment in general, which are highly correlated over time with college attainment. We tabulate this variable from the micro CPS data and include these demographically disaggregated college degree rates for ages 27 and over for both sexes.

(4) *Life expectancy* conditional on survival to each age, 55 to 79. Because mortality and morbidity at older ages tend to improve together,

45. We use second-stage final weights of the individual data to construct the participation rates, because composite weights are not available for the earlier years of our sample. In addition, we apply the seasonal factors provided by the BLS for various age groups to seasonally adjust the data; each age is assigned the seasonal factor for the smallest containing age group for which the BLS provides a seasonal factor.

46. All variables except these two unemployment rate gaps are normalized to have mean zero and variance one, in order to facilitate comparisons across coefficients.

this variable is intended to represent changes in both.⁴⁷ Higher life expectancy should increase participation by raising the level of assets necessary to finance retirement at any given age. Lower morbidity—better health at higher ages—should be associated with a lower disutility of participation.

(5) The *Social Security* “payout rate.” This is the average fraction of the Primary Insurance Amount that a person would receive if he or she were to retire at a particular age.⁴⁸ For persons below the statutory normal retirement age, a higher value implies a smaller penalty for retiring early. For older persons, a higher value of the variable implies a greater reward for delaying retirement.

(6) *Marriage and young children*. The associations between labor force participation and the presence of young children appear to vary by marital status, and vice versa. We therefore include three variables to capture this interaction, at least crudely: the percentage of women who are married with a child less than 6 years old, the percentage of women who are not married with a child less than 6 years old, and the percentage of women married without a child less than 6 years old.⁴⁹

(7) The ratio of the effective *minimum wage*, adjusted to account for state-level minimum wages that are above the federal level, relative to average hourly earnings.⁵⁰

(8) The ratio of the *median hourly wage rate for ages 16–19* to the median hourly wage rate for ages 25-plus.⁵¹ This variable is intended to reflect movements in the relative demand for teenagers, perhaps due to the factors discussed above.

(9) The ratio of summer to non-summer *school enrollment rates*, by age and sex, among teenagers. This variable represents the expansion of

47. We use sex-specific estimates of life expectancy from the Census Bureau. We include this variable for ages 55 and over for both sexes.

48. We include the payout rate for ages 62–79 for both sexes.

49. We tabulate these percentages, by age, from the micro CPS data. We include the first two of these variables for women ages 18 to 45, and the third for women ages 18 to 61. Very few women over the age of 45 have children less than 6 years old, so for the 46-plus ages the third variable acts simply as percent married.

50. Specifically, we define the ratio of the minimum wage, as a population-weighted average of federal and state minimums, to average hourly earnings. We include this variable for ages 16 to 19 for both sexes.

51. We tabulate this variable from micro CPS data at the annual frequency, and smooth it using an HP filter. We include this variable for ages 16 to 19 for both sexes.

schooling more generally since the mid-1990s, which, as noted above, has mainly occurred in the summer months.⁵²

(10) The number of *Social Security Disability Insurance* recipients, by age and sex.⁵³

In what follows, we treat the unemployment rate gap (including lags) as a measure of labor market strength or weakness and the personal bankruptcy rate variable as an indicator of household balance sheets, which are also influenced by aggregate economic conditions.⁵⁴ We treat these two variables as constituting the cyclical component of the participation rate. All of the other variables, as well as the age and cohort effects, we treat as elements of the trend in participation. As noted earlier, however, the line between cycle and trend is not always clear and bright, as several of the right-hand-side variables may be influenced by the business cycle to some extent. We will attempt to quantify the extent of this possible misattribution below.

In general, this model attempts to capture, in a parsimonious manner, many of the factors that we touched upon earlier in the paper. The age effects should capture the contribution of the aging of the population. The asymmetric unemployment rate gap should capture cyclical effects. The additional right-hand-side variables are included to account for some of the other considerations discussed previously, such as retirement (proxied by our variables for life expectancy, Social Security generosity, and age effects in general), disability (proxied by the number of SSDI recipients), trends in school enrollment (proxied by the ratio of summer to non-summer enrollment), and possibly polarization.⁵⁵

52. Using this ratio abstracts from some of the noise in quarterly enrollment rates. We tabulate this from micro CPS data. We include this variable for ages 16 to 19 for both sexes. However, we enter this variable in a different fashion from the others. (See footnote 62.)

53. By age and sex, we divide the number of SSA disability recipients in current payment status by the population for that age and sex computed from micro CPS data. We include this variable for ages 30 to 64 for both sexes. Because recipients rarely stop receiving Social Security disability payments until they “age out” or die, we essentially treat reciprocity as a predetermined variable with respect to labor force participation.

54. Later in this section, we discuss the suitability of the unemployment rate as a measure of the cycle.

55. The model likely captures some of the effects of polarization indirectly through the inclusion of right-hand-side variables that are affected by polarization or correlated with its effects, such as disability insurance take-up and the college share. However, it is unlikely that these variables capture the full impact of polarization, and if so, then the effect of labor market polarization on participation may be partly absorbed by the cohort effects, creating a sort of “endpoint” bias for more recent cohorts.

The model is estimated by least squares as a panel of 128 equations with cross-equation restrictions: one equation for each age/sex combination, with the cohort effect for a given birth cohort and sex assumed to be constant across equations.⁵⁶ The coefficients on the right-hand-side variables may vary freely across age/sex combinations except for the many cases where they are constrained to be zero, as noted above.^{57,58}

As is well known, age, cohort, and time effects are separately identified only up to a shared linear trend. We chose to restrict the time effects to be a linear combination of a fairly large set of explanatory right-hand-side variables, as we believe this maximizes the economic content of the model. The cohort and age effects are identified (with the cohort effects for 1975 normalized to zero) in our model because there are assumed to be no time effects that are not captured by the right-hand-side variables.

The 2006 Brookings paper and Fallick and Pingle (2007) argued that the cohort dimension is historically both important and meaningful and, therefore, that restricting the time dimension to elucidate the cohort dimension is a worthy choice.⁵⁹

This latter restriction is especially apt for women. As John Durand (1948, p. 123) observed, “As they grew older, each successive generation of women seems to have retained the greater propensity to be in the labor force which it developed in early adulthood, and so the higher percentages of labor force members have gradually been transmitted throughout the age groups from the late 20’s to the early 60’s.” This can be seen in figure 8, which shows the participation rates for three age groups of women (35–44, 45–54, and 55–64).⁶⁰ The horizontal axis shows the birth year for the middle age of the group. In this way, each birth cohort is vertically aligned with itself at different ages. The importance of the cohort dimension can be seen in the inflection points, which occur within each age group

56. Each age-sex-period observation is weighted based on the corresponding sample size and the value of the participation rate, under the assumption that any error in the result will be proportional to the associated log-odds transformed binomial variance for the given participation rate value.

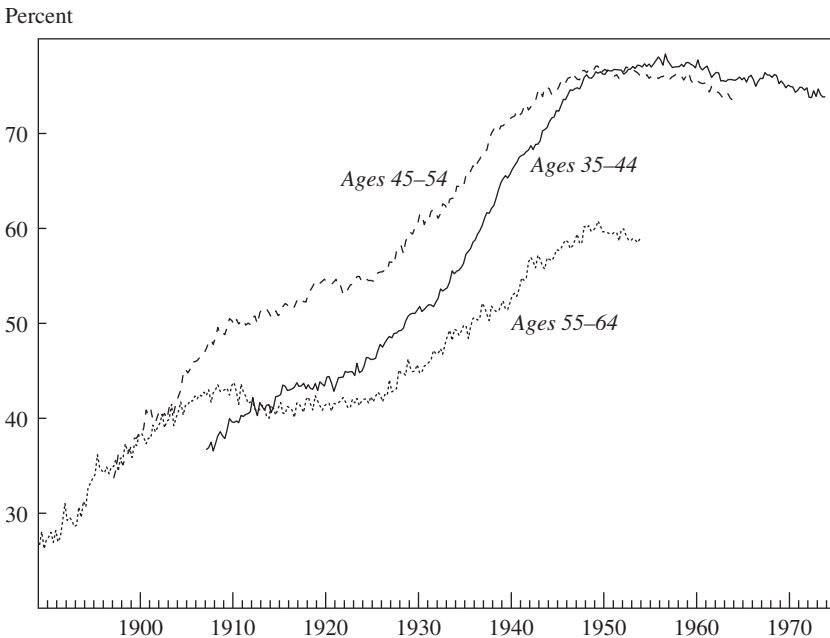
57. Because there are no constraints that cross sexes, the system is effectively estimated separately for men and for women.

58. The estimated contributions to the recent decline in participation are similar when the coefficients are constrained to be constant within various-sized age groups.

59. Numerous other authors have recognized the importance of variations by birth cohort in the context of participation as well. Some examples are Aaronson and others (2014b); Clogg (1982); Easterlin (1961); Euwals, Knoef, and Daniel van Vuuren (2011); Fernandez (2013); Goldin (2006); Lee (2014); and Vere (2007).

60. We use age groups because the data for the figure go back farther than data for individual ages are available.

Figure 8. Female Labor Force Participation Rates, Selected Age Categories, by Approximate Birth Year, 1890–1975



Source: Bureau of Labor Statistics.

at approximately the same birth cohorts. Ideally, of course, one would like to model all of the economic, technological, and social factors that led to these “greater propensit[ies] . . . developed in early adulthood,” but that has so far not proved possible.⁶¹

The apparent importance of differences in the participation rates of successive cohorts means that the evolution of the aggregate participation rate ought to be at least partly predictable. Indeed, a model with a cohort structure would have predicted the leveling off of the female participation rate that occurred around 1990, and it was one element in the prediction of further declines in participation in the earlier Brookings paper. Moreover, when it comes to the projections we show in section V, restricting the time dimension to observables allows us to make reasonable, or at least clearly defined, assumptions about how the various factors represented by the right-hand-side variables will evolve. We recognize that any time-varying

61. The pattern for men is less indicative of a cohort structure, as opposed to a time structure, for the model.

influences not captured by the included right-hand-side variables may contaminate the cohort effects, causing the projections from the model, which propagate the cohort effects through the age distribution over time, to suffer accordingly.

IV.B. Recent Cohorts

Estimates of the cohort effects for cohorts that appear for only a few years at the beginning or end of our sample period are likely to suffer from endpoint bias if estimated as part of the system. Most importantly, the cohort effects for those born between 1992 and 1998 would currently be estimated using only data since the onset of the Great Recession. Not only might this give a misleading impression of how much of the recent movements in participation are trend as opposed to cycle, but because the cohort effects are assumed to persist for a lifetime this bias could prejudice the model's predictions for future years. To mitigate this problem, we estimate the model excluding the first and last 10 cohorts from the data, for estimation of both their cohort effects and the coefficients on the right-hand-side variables. The cohort effects for these cohorts are then linearly extrapolated from the adjoining 10 cohort effects.⁶² This is one of the changes from the 2006 model, in which we included these cohort effects in the estimation but constrained them to evolve slowly.

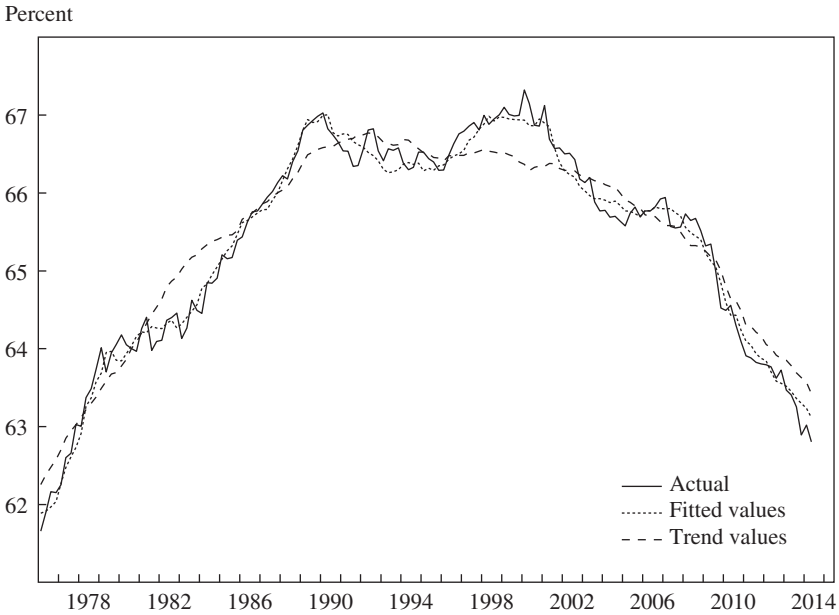
IV.C. Model Estimates

We estimated the model on quarterly data over the period 1976Q1 to 2014Q2. Figure 9 shows the estimates from the baseline model.⁶³ The solid line in the figure shows our calculated value for the actual labor force participation rate.⁶⁴ The dotted line shows the fitted values from the model. For the most part, the model tracks the actual rate reasonably closely. The

62. This approach leaves us unable to estimate in a satisfactory fashion the coefficients for the enrollment variable described above, for which the main variation occurs relatively recently. As a result, we omit the enrollment rate from the main procedure, regress the residuals for teenagers from the model on the enrollment variable in a second stage, and use the coefficients from this second-stage regression to recalculate the fitted values for teenagers.

63. As noted above, our baseline specification includes lags of the two unemployment rate gap variables at 4, 8, and 12 quarters. This choice was based on considerations of standard statistical criteria, which do not provide a definitive answer, and comparability with others' work, such as the recent report from the Council of Economic Advisers, and the panel regressions in section III. Results using no lags, fewer lags, or adding 16-quarter lags are quite similar to the baseline specification.

64. As noted above, these rates differ slightly from the published rates.

Figure 9. LFPR: Actual and Model, 1976–2014

Source: Authors' estimates using data from the U.S. Census Bureau, the Bureau of Labor Statistics, the Congressional Budget Office, the Social Security Administration, and U.S. Bankruptcy Courts.

dashed line shows the model's estimate of the trend participation rate, which, as noted above, we define as the fitted values with the unemployment gap (representing labor market slack) and the de-trended personal bankruptcy rate (representing wealth effects) set to zero. The difference between the fitted and trend values is the model's estimate of the aggregate cyclical response of the participation rate.

According to the model's estimates, the trend participation rate fell 2.1 percentage points between 2007Q4 (at which point the participation rate was a little above trend) and 2014Q2, accounting for $\frac{3}{4}$ of the total decline of 2.8 percentage points in the aggregate participation rate over this period. The contribution of cyclical factors (labor market slack plus wealth effects) to the decline is estimated to be about $\frac{1}{2}$ percentage point. In addition, as of 2014Q2, the model estimates that continued weakness in the labor market was holding the level of the participation rate down by 0.1 percentage point, with the currently low bankruptcy rate contributing a slightly larger amount, for a total cyclical response on the level of

participation of 0.3 percentage point; in addition, the participation rate in 2014Q2 was 0.3 percentage point lower than the model could explain.⁶⁵ (Of course, by construction the model cannot capture any change in the cyclical sensitivity of participation during the recent episode, a point to which we will return below.) Compared to our analysis in section III, the model estimate of the current cyclical element is about $\frac{1}{4}$ percentage point smaller than a strict read of the increase in the number of discouraged workers would suggest, and $\frac{3}{4}$ percentage point smaller than implied by the comparable lag structure in our panel data analysis. One reading of these discrepancies is that the model remains agnostic about 0.3 percentage point (the model residual) of the shortfall in participation that these earlier analyses attribute to the cycle.


The decline in the aggregate trend combines the changing age distribution of the population with disparate movements in age/sex-specific trends (which include the influence of the cohort effects). As can be seen in figure 10 the model produces reasonable trends for broad age/sex groups in the sense that the trends accord with the general patterns of participation we noted earlier. These age-specific trends combine with the changes in the age distribution to produce the overall downward aggregate trend.⁶⁶

We are comfortable attributing some of these age-specific trends to particular explanatory variables. However, some of the variables in the model are highly correlated with each other, most notably fairly monotonic variables like education levels and longevity, so the model's attributions should be interpreted with caution. Moreover, as noted above, we intended some variables to represent a broader set of related factors.

That said, looking over the longer sweep of the past four decades, the model attributes a large part of the rapid increase in prime-age women's participation rates through around 1990 to an increase in the share of women who are unmarried or without young children and to rising

65. One may be surprised that the model does not indicate a larger cyclical decline in the depths of the Great Recession than during, say, the 1990s recession. This is because the large increase in the personal bankruptcy rate during the Great Recession, and the more general deterioration in household balance sheets that this presumably indicates, worked to hold the participation rate up by inducing some affected individuals to remain in the labor force; we classify this as a cyclical response.

66. Note that the model estimates trends for single-year ages. For purposes of illustration, we aggregated these into broad age groups. Because the aggregation uses contemporaneous population shares, changes in the age distributions within groups affect the group trends. This is particularly important for the 55-and-older group. The analysis below is based on the single-year age trends.

Figure 10. Labor Force Participation Rate: Model Trends for Selected Age/Sex Groups, 1976–2014

Source: Authors' estimates using data from the U.S. Census Bureau, the Bureau of Labor Statistics, the Congressional Budget Office, the Social Security Administration, and U.S. Bankruptcy Courts.

educational attainment, counteracted to some extent by increases in disability. However, it leaves some of the increase to be accounted for by changes in cohort-specific proclivities to participate, which may reflect changes in such diverse elements as societal attitudes, workplace technologies, and reproductive technologies, to the extent that these are not already reflected in the right-hand-side variables. Meanwhile, the model attributes some of the long-standing decline in male participation in this age group to an increase in disability rolls, but it leaves the bulk of the downtrend to be “explained” by the cohort effects. The cohort effects are not much of an explanation, of course; they are more of a description. We suspect that the downtrend in cohort effects reflects diverse factors, such as increased availability of other sources of income (such as transfer programs and labor income from other members of their households) and changes in the structure of labor demand (such as deindustrialization and polarization) that were unfavorable to many prime-age men. In future work we will attempt to quantify some of these factors.

Table 3. Estimated Contributions to the Change in LFPR, 2007Q4 to 2014Q2

Percentage points

<i>Source of change</i>	<i>Contribution</i>
Age distribution	-1.3
Cohort effects	-1.7
Unemployment rate gap	-0.3
Bankruptcy rate	-0.2
Percent with college degree	+0.2
Life expectancy	+0.3
Social Security payout rate	+0.1
Marriage × young children	+0.8
Minimum wage	0.0
Teenage wage ratio	0.0
Summer enrollment ratio	0.0
Disability insurance	-0.5
Model residual	-0.2

Source: Authors' estimates using data from the U.S. Census Bureau, the Bureau of Labor Statistics, the Congressional Budget Office, the Social Security Administration, and U.S. Bankruptcy Courts.

Among persons of both sexes ages 55 and up—the range for which the option of retirement is most salient—the model attributes much of the marked increase in participation rates since the mid-1990s to a number of factors, including changes to Social Security rules and increased levels of education, but the largest contributor is the increasing life expectancy of men. The entrance into this age group of female cohorts with greater attachment to the labor force also contributed, on net, over this period. This rising contribution of cohort effects is estimated to have stopped around 2010.

Participation rates for teenagers and young adults have been falling since about 1990. The model attributes that trend decline primarily to falling cohort effects, reinforced to a small extent by increasing school enrollment among teens and counteracted to some extent by changes in fertility and marriage patterns.⁶⁷ However, the model has the greatest difficulty explaining the behavior of this group.

What of the model's interpretation of the changes since 2007? For the most part, the model attributes this decline to the same factors that have been at work over the previous two decades (table 3): A changing age distribution, falling cohort effects, rising disability rolls, and a net cyclical

67. As noted above, the model does not include the school enrollment variable for young adults.

weakening have pushed participation down more than changing patterns of marriage and fertility and increases in education and life expectancy have pushed them up.⁶⁸ The steepening of the downtrend over this period is mainly a function of the changing age distribution, as more of the large baby-boom cohort has moved farther into the age range in which the largest drop-offs in participation rates are observed, although there is also some steepening of the decline in cohort effects.

The contributions of the elements that may be described as demographic—the age distribution, marriage \times fertility, the two education variables, and life expectancy—sum to near zero over this period. However, as will become important in section V, while we can be confident about how the age distribution will roughly evolve from this point forward, whether the other demographic variables will continue to move in an offsetting direction is a more speculative matter.

Econometric studies have traditionally found that the aggregate participation rate varies little (about its trend) over the business cycle. With a glance at the movements around recessions and recoveries evident in figure 9, it should come as no surprise that our model finds a similarly small amount of cyclicity in the aggregate. By the model's estimates, a sustained 1-percentage-point increase in the unemployment rate, all else equal, currently would be expected to reduce the participation rate by something on the order of 0.2 percentage point.

As noted above, however, some of the right-hand-side variables in the model may themselves vary with the state of labor demand, and it is possible that the model does an inadequate job of attributing this indirect influence to the unemployment rate. In order to gain a sense of the potential size of any misattribution, we attempted to isolate the components of the suspect variables—bankruptcy, college attainment, marriage \times young children, teenage/adult wage ratio, summer/non-summer enrollment, and SSDI reciprocity—orthogonal to the cyclical state of the labor market by regressing them on the unemployment rate variables. We then replaced these variables in the model with the residuals from these regressions. The resulting version of the model would attribute an additional 0.3 percentage point of the decline in aggregate participation since 2007 to labor market weakness (boosting the total cyclical effect to 0.8 percentage point), primarily by increasing the estimated contribution of labor market strength in 2007.

68. Note that the contribution of the cycle to the change in the table is larger (in magnitude) than the 0.2 percentage point cited earlier, because the unemployment rate was still below the natural rate in the fourth quarter of 2007.

IV.D. The Unemployment Rate as the Model's Measure of Slack

The unemployment rate gap is an imperfect measure of the cycle for our purposes. If part of the decline in the unemployment rate is due to a cyclically depressed participation rate, the model-implied cyclical participation rate component will be biased. In particular, if the degree and persistence of cyclical weakness in the demand for labor discouraged an unusually large number of persons from remaining in or entering the labor force, then the decline in the unemployment rate could overstate the degree of labor market improvement and therefore understate the contribution of labor market slack to the decline in participation.

Given the historical patterns embodied in the model's estimates, we believe that this is probably not a major problem, although it remains an area of concern. Unfortunately, the most commonly used alternatives as cyclical indicators, such as the capacity utilization index, output gap or (detrended) nonfarm payrolls, have their own flaws, which are arguably as important as those of the unemployment rate. The capacity utilization index only measures utilization for the manufacturing sector, not the whole economy. The output gap depends on an estimate of potential output that is subject to substantial uncertainty and has been revised considerably since the Great Recession. Moreover, estimates of potential output either explicitly or implicitly depend on estimates of trend labor force participation, which is the primary quantity the model seeks to estimate. Similarly, measures of payroll growth must necessarily be detrended to account for changes in labor supply, and thus they will depend on an estimate for trend participation.

The Conference Board's survey of households, which asks respondents for their views on current labor market conditions, provides a measure that may be less subject to these criticisms. Therefore, we replaced the unemployment rate gap in the model with the difference between the percentage of households reporting (in that survey) that jobs are plentiful and the percentage reporting that jobs are hard to find.⁶⁹ Because the Conference Board measure is not available as early as is the unemployment rate, we estimated the model with a starting date of 1981. Table 4 compares the estimates using these two measures of labor market slack. The estimated contributions of the trend components (not shown) and the cycle to the decline in the participation rate are very similar to the baseline model.

69. This measure is highly correlated with a range of indicators of the labor market conditions. See Chung and others (2014). Note that the survey is unavailable at the state level, so we cannot replicate this exercise for our earlier state panel regressions.

Table 4. Estimated Contributions to the Change in LFPR (2007Q4 to 2014Q2):
Alternative Measures of Slack
(Percentage points)

<i>Source of change</i>	<i>Contribution</i>		
	<i>Unemployment gap (baseline)</i>	<i>Unemployment gap (starting 1981)</i>	<i>Conference Board (starting 1981)</i>
Labor market slack	-0.3	-0.3	-0.3
Bankruptcy rate	-0.2	-0.1	-0.1
Model residual	-0.2	-0.2	-0.3

Source: Authors' estimates using data from the U.S. Census Bureau, the Bureau of Labor Statistics, the Congressional Budget Office, the Social Security Administration, the Conference Board, and the U.S. Bankruptcy Courts.

IV.E. Model Uncertainty and Robustness

Sampling variation implies considerable uncertainty around the model's estimates of trend and cycle.⁷⁰ However, we believe that the more important issue in this case is model uncertainty. As indicated by comparing the model's results to those of the panel regressions, different approaches to the question yield different answers. However, even within the structure of the model there are consequential choices of specification that can lead to important differences in the model's characterization of the past several years. This section concentrates on two such choices.

Given the severity of the Great Recession and the drawn-out nature of the recovery, some have argued that the cyclical sensitivity of participation during this episode was greater than has historically been typical. No time-series model could reliably identify a change in the cyclical coefficients in a single episode, and there is no indication in the model's residuals of nonlinearities with respect to the severity of downturns in previous recessions.⁷¹ However, of particular concern in our model is that the presence of "excess cyclicity" in the past 7 years may influence the estimated cohort effects, especially for those cohorts with relatively few years in the data so far.

As described above, we attempted to mitigate any possible endpoint biases by estimating cohort effects only for cohorts with at least 10 years of available data and extrapolating cohort effects for the remaining cohorts.

70. The full set of coefficient estimates with associated standard error bands are graphed in the online appendix, as are confidence bands around the model's estimate of trend and cycle.

71. Of course, with only three previous recessions in our sample period, only one of which would be classified as severe, identifying nonlinearities of this sort also would be difficult.

Table 5. Estimated Contributions to the Change in LFPR (2007Q4 to 2014Q2): Varying Cohort Extrapolation and Estimation Periods
Percentage points

<i>Source of change</i>	<i>Contribution</i>		
	<i>Baseline</i>	<i>Extrapolate 15 cohorts</i>	<i>End estimation in 2007Q2</i>
Labor market slack	-0.3	-0.3	-0.7
Bankruptcy rate	-0.2	-0.1	+0.2
Trend	-2.1	-1.5	-1.5
Model residual	-0.2	-0.8	-0.7

Source: Authors' estimates using data from the U.S. Census Bureau, the Bureau of Labor Statistics, the Congressional Budget Office, the Social Security Administration, and U.S. Bankruptcy Courts.

Given the length of the current period of labor market weakness, this might not be adequate. There is a trade-off, of course, between limiting endpoint bias by reducing the number of cohorts included in the estimation and limiting the currency of the information used to estimate the model, and the proper balance is difficult to know. We explored the sensitivity of the model's estimates to this trade-off in two ways.

First, we varied the number of cohorts whose effects are extrapolated by varying the minimum number of years of data a cohort has to have available in order to be included in the estimation. It turns out that the estimated degree of cyclicity is not sensitive to varying the number of cohorts estimated. (For example, the variation in the estimated cyclical departure from trend in 2014Q2 amounts to less than 0.1 percentage point.) In contrast, the estimated trends were sensitive to this choice, mostly, although not exclusively, from variation in the estimated cohort effects. Reducing the number of extrapolated cohorts from the baseline model's total of 10 makes little difference. However, as we increase the number of cohorts extrapolated the estimated downward trend becomes less steep, and as a result more of the decline in participation over this period is left unexplained by the model. For example, the first and second columns of table 5 show the estimated contributions of trend, the two elements we identify as the cycle, and the model residual to the change between 2007 and mid-2014 in the (model-data-consistent) participation rate under the baseline model and when we extrapolate the past 15 cohorts instead of only the past 10. (Other choices can be seen in figure A3 of the online appendix.)⁷²

72. Online appendixes for papers in this volume may be accessed at the Brookings Papers webpage—<http://www.brookings.edu/about/projects/bpea>—under Past Editions.

Second, we kept the number of extrapolated cohorts constant at 10, but varied the estimation period. In this way, we can remove the Great Recession and its aftermath from the estimation completely. The third column of table 5 shows the decomposition if we end estimation in mid-2007. Reducing the number of years used in the estimation leaves the overall cyclical piece little changed, just as reducing the number of cohorts included in the estimation (that is, increasing the number of cohorts extrapolated) did. However, ending the estimation in mid-2007 causes the model's trend to fall less steeply and leaves more of the net decline in participation unexplained.^{73,74}

As with the first exercise, the model's estimate of the contribution of the cycle to the recent decline in the aggregate participation rate changes little as the end date of estimation changes. However, the composition of that contribution does change. Somewhat counterintuitively, without the Great Recession in the estimation period, the unemployment rate gap explains more of the decline, while the bankruptcy rate actually contributes positively to the change. Evidently, and not surprisingly given the relatively small and transient previous movements in the (detrended) bankruptcy rate, the large losses of the recent housing bust and financial crisis provide the main identification of these latter coefficients in the baseline estimation.

We have little to guide us to determine whether our baseline specification or one of these alternatives is more appropriate. Our baseline specification allows us to incorporate more of the recent data while enforcing a structure that we think is supported by historical experience.⁷⁵ Presumably

73. The current model estimated through 2005 produces a noticeably smaller decline in the trend since 2007 than was projected by the model in the 2006 Brookings paper. As noted above, we made several changes to the current version of the model relative to the 2006 vintage. In addition, the current exercise uses the actual values of the right-hand-side variables for the post-2005 period, whereas the 2006 paper held many of these variables constant over what was then the projection period. Of these, the most important for explaining this discrepancy appear to be a new method of extrapolation of cohort effects for the youngest cohorts and the use of actual post-2005 values for the right-hand-side variables.

74. Endpoint bias in this dimension can work in either direction. For example, if the housing boom of the mid-2000s pushed the participation rates of some groups above their trend, as might be suggested by Charles, Hurst, and Notowidigdo (2013), then ending the estimation with 2007 may bias the model's trend upward. However, as shown in figure A2 of the online appendix, the model's estimated trend falls less steeply as we move the end date of estimation earlier, regardless of whether doing so ends the estimation in a period with a stronger or weaker labor market.

75. In particular, the elimination of a greater number of cohorts from the estimation risks missing important information in the behavior of those born more than 25 years before the recession began—for example, the steep decline in the participation of young people that long preceded the recession.

as a result, the baseline leaves little of the decline in participation since 2007 unexplained, while the alternative specifications suggest that the participation rate declined more than would normally be expected in the wake of the Great Recession.⁷⁶ This distinction is important for informing how we should expect the participation rate to evolve going forward, a subject to which we now turn.

V. Future Prospects

As the unemployment rate continues to move toward its natural rate, we would expect the participation rate to move back toward its trend, which could simply entail the participation rate moving sideways or declining less quickly relative to its downward trend. For instance, if the unemployment rate were to continue to decline at its pace of recent years of about $\frac{1}{4}$ percentage point per quarter, it would reach the CBO's estimate of the long-term natural rate of unemployment in 2015Q1. Under this assumption, and given the lags in the baseline specification, the model would predict the participation rate to move roughly sideways into alignment with its trend level of about 63 percent soon thereafter.

However, as we have described above, our cross-state panel regressions and the alternative specifications of our model suggest a larger cyclical recovery—perhaps as much as 1 percentage point. Moreover, the line between being out of the labor force and being unemployed can be fluid, and it is possible that some people whom we might consider to have dropped out of the labor force for structural, rather than cyclical, reasons could re-enter it if appropriate opportunities arose. Indeed, the participation rate was above the model's estimate of the trend just prior to the financial crisis, when the unemployment rate was low. Therefore, in this section we discuss some ways to think about likely future outcomes, differentiating between developments over the next few years, when the size of any additional cyclical increase in participation is most relevant, and the longer run, when model estimates of the trend decline in participation are more relevant.

V.A. Changes in Labor Force Participation over the Next Few Years

Our first step in exploring the scope for cyclical recovery is to examine the size of the pool of potential labor force entrants and to judge whether

76. The report of the Council of Economic Advisers (2014) attributes 1 percentage point of the decline in the participation rate between the end of 2007 and mid-2014 to unexplained factors, similar to our alternative in which we estimate fewer cohorts.

their entry into the labor force could have a meaningful impact on the aggregate participation rate. Nonparticipants who report that they want a job would seem to be the most salient pool of potential labor force entrants, and indeed (as shown in the top panel of figure 11) more than 40 percent of such individuals who report that they want a job in one month flow back into the labor force 12 months later.⁷⁷ Moreover, based on our analysis (not shown) of data from the CPS March supplement, these potential workers are also more likely to have spent time working in the recent past than nonparticipants who report not wanting a job,⁷⁸ and there is some evidence that they may exert some downward pressure on wages.⁷⁹ As we noted above, however, the size of this group is relatively small and would likely provide only a small boost to the participation rate going forward, on the order of ½ percentage point.

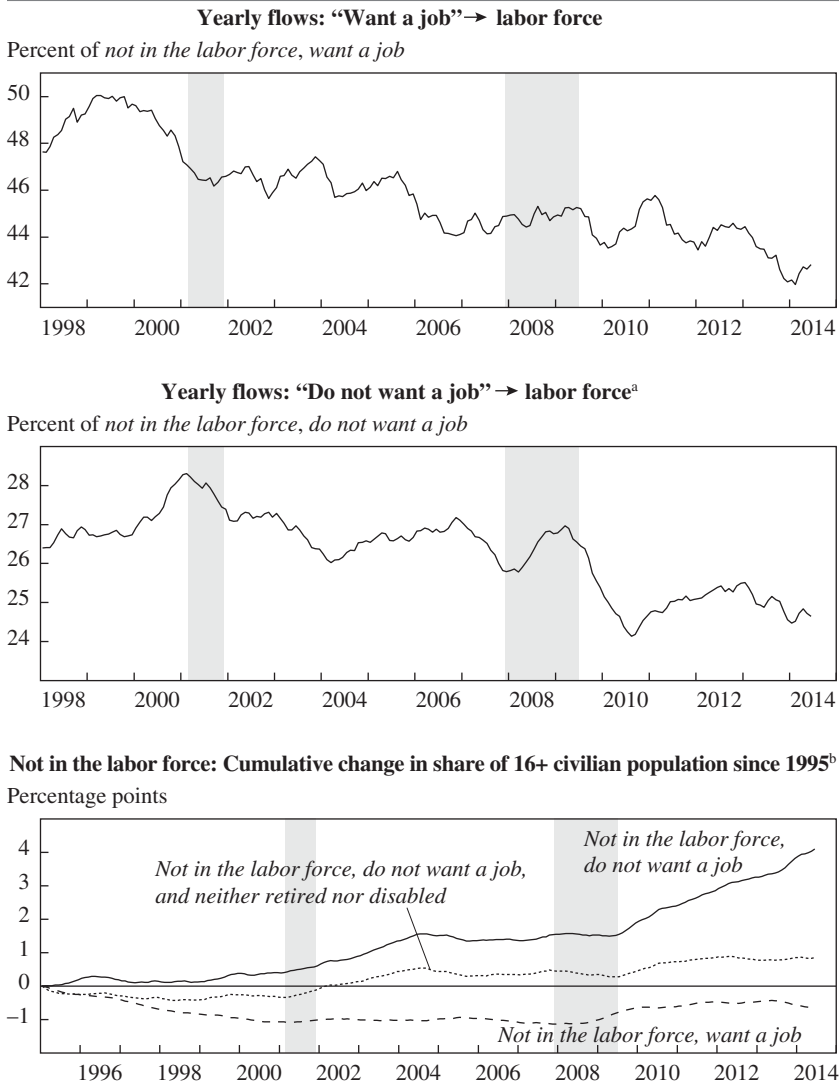
Nonparticipants who report not wanting a job are more numerous, but the evidence suggests they are less likely to re-enter the labor force. As we have noted, the increase in the ranks of retirees and the disabled in recent years appears to be largely structural, and such individuals have low transition rates into the labor force. That said, there may be nonparticipants who are currently engaged in other activities (such as attending school or raising children) and report they do not want a job but who could be drawn back into the labor force as the economy continues to strengthen. For example, the transition rate for a nonretired, nondisabled person who is out of the labor force and does not want a job (figure 11, middle panel) is above 20 percent—not as high as for those who want a job, but not trivial either.

77. These 12-month labor force flows as well as those in the earlier discussion of retirement and disability relate to the flows between five states: in the labor force, not in the labor force but want a job, retired, disabled, and all other not in the labor force and do not want a job. They are calculated by the authors from the Current Population Survey Longitudinal Population Database (see Nekarda 2009).

78. Specifically, using the March Annual Social and Economic Supplement to the CPS we identified those reporting themselves as marginally attached in the basic March survey and then examined their work history during the prior calendar year using the Supplement. This analysis shows that in recent years about 30 percent of those reporting themselves as wanting a job in March worked at least one week in the prior year, compared to fewer than 10 percent of those who do not want a job. Those reporting themselves as wanting a job were also more likely to have looked for work in the prior year (over 25 percent in recent years, compared to about 2 percent of those not wanting a job). That said, these individuals were still less likely to have looked for work or worked in the prior year than those who reported themselves unemployed at the time of the March survey (about 57 percent of whom engaged in work and or search).

79. Indeed, cross-state evidence suggests that a 1-percentage-point increase in the share of the population that wants a job holds down wage growth by 0.7 percentage point (see Smith 2014).

Figure 11. Not in the Labor Force (NLF), by Want-Job Status, 1998–2014



Source: Authors’ calculations using data from the Current Population Survey Longitudinal Population Database (see Nekarda 2009) and from the U.S. Census Bureau and the Bureau of Labor Statistics (Current Population Survey).

a. The flows in the middle panel refer to those who are not in the labor force, do not want a job, and are not retired or disabled.

b. 12-month moving averages.

However, there has not been a particularly large increase in the share of such individuals in recent years. This is illustrated in the bottom panel of figure 11, which plots the cumulative change since 1995 in the share of the 16-and-over population reporting themselves out of the labor force and not wanting a job. As shown by the solid line, this series turns up sharply after the end of the Great Recession. However, once we eliminate retirees and those reporting themselves as disabled, groups for whom (as we argued above) the decline in participation was largely structural,⁸⁰ the fraction not in the labor force and not wanting a job (the dotted line) has moved up only slightly since the recession, accounting for perhaps a ½ percentage point decrease in the participation rate.⁸¹

Another way to consider the potential scope for a rebound in participation is to examine whether inflows to the labor force are atypically low for this point in the cycle.⁸² In figure 12, we plot the monthly transition probabilities from out of the labor force into employment or unemployment (top and middle panels) as well as into the labor force as a whole (bottom panel) against the monthly aggregate unemployment rate, along with estimated regression lines. It turns out that the nearly acyclical nature of the participation rate is the product of two opposing factors. The monthly rate at which individuals enter the labor force for employment (NE flow) declines as the unemployment rate increases (top panel), no doubt as a result of reduced employment opportunities.⁸³ But at the same time, monthly transitions from outside of the labor force into unemployment (NU flow) increase (middle panel). It may be that the increase in NU flows is also related to poor labor market demand, as individuals who ordinarily would have moved from out of the labor force into a job (such as recent graduates) endure a spell of unemployment when opportunities

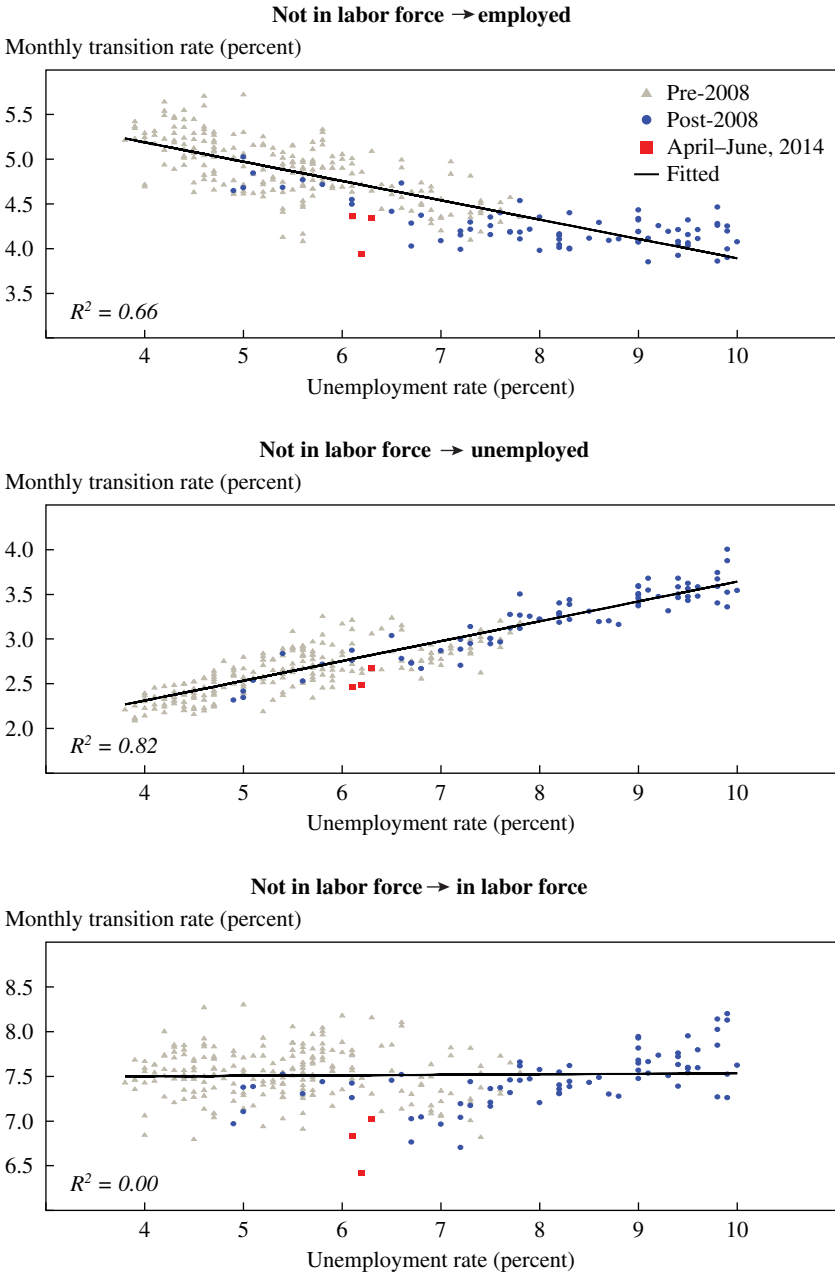
80. Of course, it is possible that there was something unique about the current episode and that these persons will return to the labor force in greater numbers than in the past. For instance, if many more workers below age 62 had retired than is typical in less severe recessions, we might expect larger flows from retirement into the labor force, since these workers appear to have greater labor force attachment. However, as shown earlier, retirement rates for those under the normal retirement age have not increased since 2007.

81. This is relative to a counterfactual that the nonretired, nondisabled share of the population reporting themselves as not wanting a job would have been flat in the absence of the recession. In fact, this share has been rising over time (Barnichon and Figura 2013), suggesting that this may be an upper bound on the potential pool of participants.

82. This analysis uses the monthly flows published by the BLS, which have been seasonally adjusted and raked to be consistent with the published levels.

83. A similar chart has been plotted by Matthew O'Brien in *The Atlantic*. See "The Fed Absolutely Shouldn't Give Up on the Long-Term Unemployed," March 12, 2014.

Figure 12. Monthly Flows into the Labor Force, February 1990–June 2014^a



Source: Bureau of Labor Statistics.
 a. Each dot represents a monthly observation.

are poor.⁸⁴ Whatever the explanation, combining the two flows results in a pattern of overall flows from not-in-the-labor-force into the labor force that is indeed acyclical (bottom panel), which is consistent with the findings of our model and would suggest the possibility of, at most, a modest inflow of workers into the labor force as the unemployment rate falls. That said, the most recent data points (the squares in figure 12) are somewhat low relative to the unemployment rate in the second quarter of 2014 (given the usual historical relationship between the series), providing some supporting evidence for a larger-than-usual cyclical rebound.

We see the greatest uncertainty for the evolution of the aggregate participation rate as stemming from the evolution of participation rates for young people. In this regard there are two questions: (i) Will future cohorts have similarly low participation when they are ages 16 to 24? and (ii) Will the low participation among today's young people persist throughout their lives?

With respect to the first question, we have argued that the recent declines in participation for teenagers and young adults appear to be a continuation of long-standing trends. So while it seems possible that enrollment rates could move down as the economy improves, which would likely provide a small boost to the aggregate participation rate, and it is also possible that (conditional on enrollment status) participation rates could rise a touch as the labor market tightens, for the most part we continue to expect young people to work less than in previous generations.

With respect to the question of whether the low participation rates of young people will translate into relatively low lifetime participation for these cohorts, our baseline assumption is that they will. Of course, over time the participation of recent cohorts will rise, in keeping with the typical age effects. Moreover, all else equal, the greater educational attainment of recent cohorts should translate into relatively higher participation rates in their adult years, as could further changes in patterns of marriage and fertility. We explore alternative assumptions in the next section.

V.B. Changes in Labor Force Participation over the Next Decade: Model Projections

Beyond any cyclical effects, which are likely to be salient only over the next year or so, the participation rate in coming years will be determined

84. A recent paper by Elsby, Hobijn, and Şahin (2013) finds the countercyclical nature of the NU flows hard to explain; according to those authors, classification errors may only partly explain it, and other commonly cited channels, such as the added worker effect, don't seem to be supported by the data.

Table 6. Alternative Projections of LFPR (2014Q2 to 2024Q2)

	<i>Model</i>		
	<i>(1)</i> <i>Equation 1</i>	<i>(2)</i> <i>Holding incoming cohort effects and most variables constant^a</i>	<i>(3)</i> <i>Holding incoming cohort effects constant and extrapolating all other variables^a</i>
2014Q2	62.8	63.1	63.1
2015Q2	62.6	62.9	63.0
2016Q2	62.3	62.5	62.7
2017Q2	62.0	62.0	62.4
2018Q2	61.8	61.5	62.1
2019Q2	61.6	61.1	61.8
2020Q2	61.3	60.6	61.5
2021Q2	61.0	60.2	61.3
2022Q2	60.7	59.8	61.0
2023Q2	60.4	59.3	60.7
2024Q2	60.2	58.9	60.4

Source: Authors' estimates.

a. Holding cohort effects for incoming cohorts constant at the last estimated value, that is, for persons born in 1998.

by its trend. We therefore turn now to the model's forecast of the aggregate participation rate.

We begin by calculating the implications of projected changes in the age distribution alone using equation 1, holding the age-specific participation rates constant at their 2014Q2 values. We use the "middle" projections from the Census Bureau to project the evolution of the age distribution from 2014Q2 on. Column 1 of table 6 shows the result of this exercise. The changing age distribution alone would be expected to lower the aggregate participation rate a further 2.6 percentage points over the next 10 years.

Next, we use the baseline model to project the aggregate participation rate, allowing the age-specific trend participation rates to continue to evolve. Following the model, the evolution of those age-specific trends depends on the paths of the cohort effects for future incoming cohorts and of the various explanatory variables. Of course, any assumptions for these paths are highly speculative. Nevertheless, they can provide a sense of the range of reasonable projections.

We examine two scenarios. In both, we assume that the unemployment rate falls at its recent pace by about $\frac{1}{4}$ percentage point per quarter until it reaches the CBO's estimate of the long-term natural rate of unemployment in 2015Q1; we hold the unemployment rate gap at zero thereafter. We project the bankruptcy rate from a linear regression on the unemployment rate gap. In both scenarios, we set the cohort effects for incoming cohorts to the value estimated for the most recent (as of 2014Q2) cohort of 16-year-olds (that is, no further declines), extrapolate a continued increase in life expectancies, and allow the Social Security payout rate to evolve according to current law. However, the treatment of the other variables differs between the two scenarios.⁸⁵

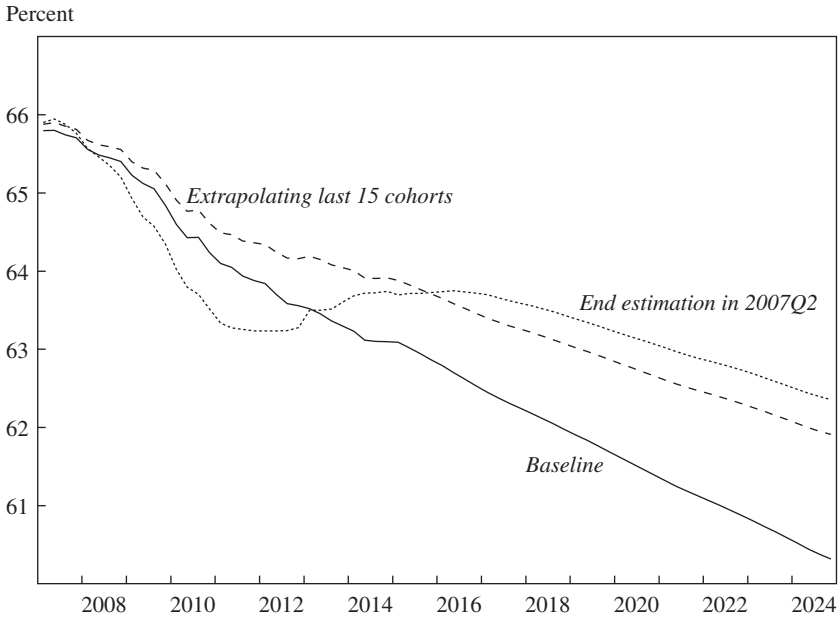
The first scenario, shown in column 2, attempts to isolate the implications of the aging of the population by holding the right-hand-side variables not mentioned above constant at their last observed levels. This scenario differs from the calculation using equation 1 along several dimensions. The most important difference is that, although it holds the cohort effects for future incoming cohorts constant at the level of the most recent incoming cohort, it allows the existing cohort effects to continue to move through the age distribution as cohorts continue to age. It is this assumption that is mainly responsible for the steeper downward trajectory of the participation rate: Because the younger cohorts have lower estimated cohort effects than the older ones, they pull down the within-age-group participation rates as they age.

In the second scenario we (linearly) extrapolate forward each of the variables held constant in the first scenario. The result of this exercise is shown in column 3 of table 6. In this case, the decline in the aggregate participation rate is similar to that in column 1, as the factors mentioned earlier—notably increasing longevity and educational attainment and changes in marriage and fertility patterns—continue, offsetting the propagation of the cohort effects.

Beginning from the scenario in column 3, we also imitate the earlier robustness exercises by varying either the number of cohorts included in the estimation or the estimation period. For simplicity, figure 13 shows the baseline projection (omitting the most recent 10 cohorts from the estimation and estimating the model through 2014Q2) extrapolating the right-hand-side variables as in column 3 of table 6, along with the two alternatives featured in section IV: one in which the estimation omits the

85. The age effects are, by construction, constant over time in all of the scenarios.

Figure 13. Alternative Model Projections for the Labor Force Participation Rate, 2007–24



Source: Authors' estimates.

most recent 15 cohorts and one in which the estimation ends in 2007Q2. By 2024, the projected participation rates from the two alternatives lie $2\frac{1}{2}$ and 2 percentage points, respectively, above the baseline projection.⁸⁶

Table 7 compares the projection for the annual average participation rate based on our second scenario (column 3 above) to projections developed by several government or international agencies. In 2014, our model's projection of the labor force participation rate, at 63.1 percent, is similar to that of the other projections. However, over the next 8 years, we project the participation rate to decline $2\frac{1}{4}$ percentage points to 61 percent, a steeper decline than that projected by the BLS, CBO, or SSA. At the same time, the differences in these projections are not as striking as they were in the 2006 Brookings paper, in which our 10-year-ahead projection for 2015 was $2\frac{1}{2}$ percentage points below that of the next lowest forecast.

86. As noted in section IV, ending the estimation period before the Great Recession robs the bankruptcy variable of influence. That is why the fitted values from that exercise fall below the baseline until 2012.

Table 7. Comparisons of LFPR 2014–2022: Authors' and Others' Models (Percent)

<i>Year</i>	<i>Authors' model^a</i>	<i>Congressional Budget Office^a</i>	<i>Bureau of Labor Statistics^b</i>	<i>Social Security Administration^a</i>	<i>International Monetary Fund^a</i>
2014	63.1	62.9	63.3	63.1	63.0
2015	63.0	62.7	63.1	63.2	63.0
2016	62.7	62.5	63.0	63.2	62.9
2017	62.3	62.4	62.7	63.2	62.8
2018	62.1	62.2	62.5	63.3	62.6
2019	61.8	62.0	62.3	63.3	62.3
2020	61.5	61.8	62.0	63.3	—
2021	61.2	61.5	61.8	63.1	—
2022	61.0	61.3	61.6	62.9	—

Sources: Authors' estimates; Congressional Budget Office (2014); Toossi (2013); Social Security Administration (2014, unpublished data); International Monetary Fund (2014).

a. Authors' projections are for the annual average participation rate, as are those from the CBO, the SSA, and the International Monetary Fund.

b. BLS projections are for the annual average trend participation rate.

VI. Conclusions

The evidence we present in this paper suggests that most of the steep decline in the labor force participation rate since 2007 owes to ongoing structural influences that are pushing down the participation rate rather than to a pronounced cyclical weakness related to potential job seekers' discouragement about the weak state of the labor market. Most prominently, the ongoing aging of the baby-boom generation into ages when attachment to the labor force is traditionally lower can, by itself, account for nearly half of the decline. In addition, estimates from our model, as well as the supplementary evidence on which we report, show persistent declines in participation rates for some specific age/sex categories, including youth and adult men. These persistent declines appear to have their roots in longer-run changes in the labor market that predate the financial crisis by a decade or more. And while increasing longevity and better health status, coupled with changes in Social Security rules and increased educational attainment, have all contributed to an ongoing rise in the participation rates of older individuals, those increases have not been large enough to greatly offset the various downward influences on the aggregate participation rate.

That is not to say that all of the decline in labor force participation reflects structural influences. Our cohort-based model suggests that cyclical weakness was depressing the participation rate by about $\frac{1}{4}$ percentage point in 2014Q2, and alternative specifications of the model yield a larger

unexplained portion of the decline that may represent unusually large cyclical. Evidence from cross-state regressions and our examination of the potential labor force pool represented by nonparticipants suggest that the contribution of cyclical weakness could be higher, even as much as 1 percentage point. Estimating the cyclical of the participation rate is a difficult econometric problem, and it is not surprising that different procedures yield different results.

Looking ahead, demographics will likely continue to play a prominent role in determining the future path of the aggregate labor force participation rate. The youngest members of the baby-boom generation are still in their early fifties, so the effects of population aging will continue to put downward pressure on the participation rate for some time. Indeed, according to our estimates, the continued aging of the population alone will subtract 2½ percentage points from the aggregate participation rate over the next 10 years. And the overall downtrend could be even larger if some of the negative trends evident for particular age-sex groups persist.

If one accepts our conclusion that the aggregate participation will likely decline further over the next decade, it is well worth considering what implications this could have for broader macroeconomic developments. The first-order effect is that—holding trends in population growth (including migration), average hours worked, and productivity fixed—the nearly 2¼-percentage-point decline in the aggregate participation rate, which we project over the next decade, will hold down trend output growth by a little less than ½ percentage point per year through the end of the decade. Another implication is that, as the growth in the labor force slows, the “break-even” level of monthly job gains required to hold the unemployment rate unchanged month-to-month will be lower than in decades past. By our calculations, over the next decade somewhere between 50,000 and 75,000 jobs per month will be needed to maintain an unchanged unemployment rate, far fewer than the amount needed in the 1990s.

Of course, much uncertainty attends these projections. While we can be reasonably sure that the domestic population will age according to Census projections, the future pace of immigration will undoubtedly influence the age distribution of the population, as new immigrants are more likely to be in their prime working years than is the general population. Moreover, future trends in participation for specific demographic groups are difficult to predict. Our model has had little success in accurately capturing changes in teenage participation rates, and given the opposing effects of increased school enrollment and polarization, future changes in participation for currently younger cohorts seem especially uncertain. Similarly, while a further

uptrend in the participation rates for older individuals seems likely, the pace of that uptrend is difficult to predict.

Moreover, our analysis does not account for general equilibrium feedbacks or future labor market policies that could mitigate future declines in participation. For instance, as the downward trend in participation restrains the growth in the labor force, firms may react by increasing real wage rates or otherwise making work more attractive—say, by making work arrangements more flexible or increasing nonmonetary compensation. If these developments come to pass and working becomes a more attractive option, those with traditionally low labor force attachment may enter the labor force in greater numbers. One optimistic scenario is that as baby boomers continue to retire, job vacancies may rise in sufficient numbers to mitigate some of the secular downtrend in participation for younger adults and less-educated workers.

In the end, however, we see further declines in the aggregate labor force participation rate as the most likely outcome. Ongoing improvements in labor market conditions may stem that decline temporarily as discouraged workers are pulled back into the job market, and indeed it would not be surprising if the participation rate moved above its trend for a time were GDP to exceed its potential level, as it appeared to do just prior to the financial crisis. Over the longer term, however, the downward influences on the aggregate labor force participation rate will likely dominate, restraining trend growth in the aggregate labor force and in the growth rate of GDP.

ACKNOWLEDGMENTS We thank Robert Hall, Jim Stock, and the editors for helpful comments on an earlier draft, and Erik Larsson, Dennis Mawhirter, and Tyler Petersen for excellent research assistance. The authors also benefited from conversations and collaboration with Andrew Figura, Charles Fleischman, Chris Nekarda, and David Ratner. The views expressed in this paper do not necessarily reflect the views of the Board of Governors of the Federal Reserve System, the Federal Reserve Bank of Cleveland, or other members of their staffs. As staff members of governmental and quasi-governmental institutions, we can also affirm that none of us has any relevant material or financial interest to declare regarding the content of this paper.

References

- Aaronson, Stephanie, Tomaz Cajner, Bruce Fallick, Felix Galbis-Reig, Christopher L. Smith, and William L. Wascher. 2014a. "Labor Force Participation: Recent Developments and Future Prospects." Finance and Economics Discussion Series Paper no. 2014-64. Washington: Board of Governors of the Federal Reserve System.
- Aaronson, Stephanie, Bruce Fallick, Andrew Figura, Jonathan Pingle, and William L. Wascher. 2006. "The Recent Decline in Labor Force Participation and Its Implications for Potential Labor Supply." *Brookings Papers on Economic Activity*, Spring.
- Aaronson, Daniel, Luojia Hu, Arian Seifoddini, and Daniel G. Sullivan. 2014b. "Declining Labor Force Participation and Its Implications for Unemployment and Employment Growth." *Economic Perspectives* Q4: 100–138 (Federal Reserve Bank of Chicago).
- Aaronson, Daniel, Kyung-Hong Park, and Daniel Sullivan. 2006. "The Decline in Teen Labor Force Participation." *Economic Perspectives* Q1: 2–18 (Federal Reserve Bank of Chicago).
- Autor, David H. 2010. "The Polarization of Job Opportunities in the U.S. Labor Market: Implications for Employment and Earnings." Policy Brief. Washington: Center for American Progress and the Hamilton Project.
- . 2011. "The Unsustainable Rise of the Disability Rolls in the United States: Causes, Consequences, and Policy Options." Working Paper no. 17697. Cambridge, Mass.: National Bureau of Economic Research.
- Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "Untangling Trade and Technology: Evidence from Local Labor Markets." Working Paper no. 18938. Cambridge, Mass.: National Bureau of Economic Research.
- Autor, David H., and Mark G. Duggan. 2003. "The Rise in the Disability Rolls and the Decline in Unemployment." *Quarterly Journal of Economics* 118, no. 1.
- Balleer, Almut, Ramon Gomez-Salvador, and Jarkko Turunen. 2009. "Labour Force Participation in the Euro Area: A Cohort-Based Analysis." Working Paper no. 1049. European Central Bank.
- . 2014. "Labour Force Participation across Europe: A Cohort Based Analysis." *Empirical Economics* 46, no. 4: 1385–1415.
- Barnes, Michelle L., Fabià Gumbau-Brisa, and Giovanni P. Olivei. 2013. "Cyclical versus Secular: Decomposing the Recent Decline in U.S. Labor Force Participation." Public Policy Brief no. 13-2. Federal Reserve Bank of Boston.
- Barnichon, Regis, and Andrew Figura. 2013. "Declining Labor Force Attachment and Downward Trends in Unemployment and Participation." Working Paper no. 2013-88, Finance and Economics Discussion Series. Washington: Board of Governors of the Federal Reserve System.
- Benito, Andrew, and Philip Bunn. 2011. "Understanding Labour Force Participation in the United Kingdom." *Bank of England Quarterly Bulletin* Q1: pp. 36–42.

- Blau, David M., and Ryan M. Goodstein. 2010. "Can Social Security Explain Trends in Labor Force Participation of Older Men in the United States?" *Journal of Human Resources* 45, no. 2.
- Bosworth, Barry P., and Gary Burtless. 2010. "Recessions, Wealth Destruction, and the Timing of Retirement." Working Paper no. 2010-22. Center for Retirement Research, Boston College.
- Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo. 2013. "Manufacturing Decline, Housing Booms, and Non-Employment." Working Paper no. 18949. Cambridge, Mass.: National Bureau of Economic Research.
- Chung, Hess T., Bruce Fallick, Christopher J. Nekarda, and David D. Ratner. 2014. "Assessing the Change in Labor Market Conditions." Working Paper no. 2014-109, Finance and Economics Discussion Series. Washington: Board of Governors of the Federal Reserve System.
- Clark, Kim B., and Lawrence H. Summers. 1982. "Labour Force Participation: Timing and Persistence." *Review of Economic Studies* 49, no. 2: 825–44.
- Clogg, Clifford C. 1982. "Cohort Analysis of Recent Trends in Labor Force Participation." *Demography* 19, no. 4: 459–79.
- Congressional Budget Office. 2014. "An Update to the Budget and Economic Outlook: 2014 to 2024." Washington. https://www.cbo.gov/sites/default/files/45653-OutlookUpdate_2014_Aug.pdf
- Council of Economic Advisers. 2014. "The Labor Force Participation Rate since 2007: Causes and Policy Implications." Washington. http://www.whitehouse.gov/sites/default/files/docs/labor_force_participation_report.pdf
- Duggan, Mark G., and Scott A. Imberman. 2009. "Why Are the Disability Rolls Skyrocketing? The Contribution of Population Characteristics, Economic Conditions, and Program Generosity." In *Health at Older Ages: The Causes and Consequences of Declining Disability among the Elderly*, edited by David M. Cutler and David A. Wise. Cambridge, Mass.: National Bureau of Economic Research.
- Durand, John D. 1948. "The Labor Force in the United States 1890–1960." New York: Social Science Research Council.
- Duval-Hernández, Robert, and Pedro Orraca Romano. 2009. "A Cohort Analysis of Labor Participation in Mexico, 1987–2009." Discussion Paper no. 4371. Institute for the Study of Labor (IZA).
- Easterlin, Richard A. 1961. "The American Baby Boom in Historical Perspective." *American Economic Review* 51: 869–911.
- Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2013. "On the Importance of the Participation Margin for Market Fluctuations." Working Paper no. 2013-05. Federal Reserve Bank of San Francisco.
- Erceg, Christopher J., and Andrew T. Levin. 2013. "Labor Force Participation and Monetary Policy in the Wake of the Great Recession." Working Paper no. 13/245. Washington: International Monetary Fund.

- Euwals, Rob, Marike Knoef, and Daniel van Vuuren. 2011. "The Trend in Female Labour Force Participation: What Can Be Expected for the Future?" *Empirical Economics* 40, no. 3: 729–53.
- Fallick, Bruce, and Jonathan Pingle. 2007. "A Cohort-Based Model of Labor Force Participation." Working Paper no. 2007-9, Finance and Economics Discussion Series. Washington: Board of Governors of the Federal Reserve System.
- Fernandez, Raquel. 2013. "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century." *American Economic Review* 103, no. 1: 472–500.
- Foote, Christopher L., and Richard W. Ryan. 2014. "Labor-Market Polarization over the Business Cycle." In *Macroeconomics Annual 2014*, vol. 29, edited by Jonathan Parker and Michael Woodford. Cambridge, Mass.: National Bureau of Economic Research.
- Fujita, Shigeru. 2014. "On the Causes of Declines in the Labor Force Participation Rate." Research Rap Special Report, Federal Reserve Bank of Philadelphia.
- Goldin, Claudia. 2006. "The Quiet Revolution That Transformed Women's Employment, Education and Family." *American Economic Review Papers and Proceedings* 96, no. 2: 1–21.
- Gorodnichenko, Yuriy, Jae Song, and Dmitriy Stolyarov. 2013. "Macroeconomic Determinants of Retirement Timing." Working Paper no. 19638. Cambridge, Mass.: National Bureau of Economic Research.
- Hall, Robert. 2014. "Quantifying the Lasting Harm to the U.S. Economy from the Financial Crisis." In *Macroeconomics Annual 2014*, vol. 29, edited by Jonathan Parker and Michael Woodford. Cambridge, Mass.: National Bureau of Economic Research.
- Helman, Ruth, Nevin Adams, Craig Copeland, Jack VanDerhei. 2014. "The 2014 Retirement Confidence Survey: Confidence Rebounds—for Those with Retirement Plans." Issue Brief no. 397. Washington: Employment Benefit Research Institute.
- Hotchkiss, Julie L., and Fernando Rios-Avila. 2013. "Identifying Factors behind the Decline in the U.S. Labor Force Participation Rate." *Business and Economic Research* (Macrothink Institute) 3, no. 1: 257–75.
- International Monetary Fund. 2014. "United States: Selected Issues." Country Report no. 14/222. Washington.
- Jaimovich, Nir, and Henry E. Siu. 2012. "The Trend Is the Cycle: Job Polarization and Jobless Recoveries." Working Paper no. 18334. Cambridge, Mass.: National Bureau of Economic Research.
- Juhn, Chinhui. 1992. "Decline of Male Labor Market Participation: The Role of Declining Market Opportunities." *Quarterly Journal of Economics* 107, no. 1: 79–121.
- Kahn, Lisa. 2010. "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy." *Labour Economics* 17, no. 2: 303–16.
- Kawata, Hiroshi, and Saori Naganuma. 2010. "Labor Force Participation Rate in Japan." *Bank of Japan Review* 2010-E-7 (December): 1–9.

- Kudlyak, Marianna. 2013. "A Cohort Model of Labor Force Participation." *Economic Quarterly* 99, no. 1: 25–43 (Federal Reserve Bank of Richmond).
- Lee, Jin Young. 2014. "The Plateau in U.S. Women's Labor Force Participation: A Cohort Analysis" *Industrial Relations* 53, no. 1: 46–71.
- Mastrobuoni, Giovanni. 2009. "Labor Supply Effects of the Recent Social Security Benefit Cuts: Empirical Estimates Using Cohort Discontinuities." *Journal of Public Economics* 93, no. 11–12.
- Morisi, Teresa K. 2008. "Youth Enrollment and Employment during the School Year." *Monthly Labor Review*, February: 51–63.
- . 2010. "The Early 2000s: A Period of Declining Teen Summer Employment Rates." *Monthly Labor Review*, May: 23–35.
- Nekarda, Christopher J. 2009. "A Longitudinal Analysis of the Current Population Survey: Assessing the Cyclical Bias of Geographic Mobility." Working Paper. Published online at Chrisnekarda.com.
- O'Brien, Matthew. 2014. "The Fed Absolutely Shouldn't Give Up on the Long-Term Unemployed." *The Atlantic*, March.
- Ramey, Gary, and Valerie A. Ramey. 2010. "The Rug Rat Race." *Brookings Papers on Economic Activity*, Spring: 129–99.
- Reifschneider, David, William L. Wascher, and David W. Wilcox. 2013. "Aggregate Supply in the United States: Recent Developments and Implications for the Conduct of Monetary Policy." Working Paper no. 2013-77, Finance and Economics Discussion Series. Washington: Board of Governors of the Federal Reserve System.
- Smith, Christopher L. 2011. "Polarization, Immigration, Education: What's Behind the Dramatic Decline in Youth Employment?" Working Paper no. 2011-41, Finance and Economics Discussion Series. Washington: Board of Governors of the Federal Reserve System.
- . 2012. "The Impact of Low-Skilled Immigration on the Youth Labor Market." *Journal of Labor Economics* 30, no. 1: 55–89.
- . 2014. "The Effect of Labor Slack on Wages: Evidence from State-Level Regressions." *FEDS Notes* 2014-06-02. Washington: Board of Governors of the Federal Reserve System.
- Toossi, Mitra. 2011. "A Behavioral Model for Projecting the Labor Force Participation Rate." *Monthly Labor Review*, May: 25–42.
- . 2013. "Labor Force Projections to 2022: The Labor Force Participation Rate Continues to Fall." *Monthly Labor Review*, December. Washington: Bureau of Labor Statistics.
- Van Zandweghe, Willem. 2012. "Interpreting the Recent Decline in Labor Force Participation." *Economic Review* Q1: 5–34 (Federal Reserve Bank of Kansas City).
- Vere, James P. 2007. "'Having It All' No Longer: Fertility, Female Labor Supply, and the New Life Choices of Generation X." *Demography* 44, no. 4: 821–28.

Comments and Discussion

COMMENT BY

ROBERT E. HALL The substantial decline in labor-force participation in recent years has raised the important question: How much of this decline is the result of the slack labor market from the Great Recession, and how much comes from other, structural forces? As the unemployment rate has returned to normal, a concern has developed that some of the people now classified as out of the labor force are, effectively, unemployed, but they are not included in the standard unemployment count because they do not satisfy its fairly exacting standards for classifying people as unemployed. Stephanie Aaronson and the other five authors of the paper under discussion tackle this question from a variety of angles. They start by looking directly at a number of factors that might have altered the structural participation rate, then perform econometric studies of the correlation of participation with changes in labor-market slackness, measured by the unemployment rate.

To better understand some of the issues the paper considers, it is useful to have a model of the interplay of labor-market slackness and participation. The observed variables are unemployment, u_t , and the participation rate, p_t . The goal is to model their behavior in terms of two unobserved factors: job labor-market slack, s_t , and structural participation, x_t . Policy-makers are interested in inferring slack (s_t) for two reasons: (i) it measures underutilization of labor, and (ii) it drives wage inflation, according to a widely held view.

The model has two equations. The first says that structural shifts raise participation and labor-market slack lowers it:

$$p_t = x_t - \gamma s_t,$$

and the second says that unemployment is positively related to slack and negatively related to structural participation, x_t :

$$u_t = s_t - \alpha x_t.$$

In both equations, the coefficient of one of the latent factors is normalized at one, reflecting the fact that the choice of units for those factors is arbitrary.

The second equation allows for hidden unemployment. Today, there is a suspicion that there is hidden unemployment because of lingering effects of the crisis. If α is positive and structural participation is low, unemployment is high even if slackness is normal. Many observers have used this model implicitly when suggesting that lower unemployment is not signaling a tighter labor market because discouraged unemployed people are exiting the labor market.

The solution of the two equations untangles the two feedbacks and extracts s_t from the data:

$$s_t = \frac{u_t + \alpha p_t}{1 - \gamma\alpha}.$$

Not surprisingly, learning s_t requires knowledge of the feedback coefficients α and γ . This model is a classic two-equation simultaneous system, like a model of supply and demand. The latent variables s_t and x_t are the disturbances. Without additional strong assumptions, estimating α and γ is impossible, and so is inferring the latent variables.

The authors' paper overcomes this problem in one respect—it brings in additional information about the structural component of nonparticipation, x_t . But its main econometric exercises—the state panel study and the participation model based on micro data—do not consider the problem. Rather, the paper makes the implicit assumptions needed to recover γ , the effect of slackness on participation, from the regression of participation on unemployment. That regression coefficient is

$$\frac{\text{Cov}(p, u)}{V(u)} = \frac{\gamma\sigma_s^2 - \alpha\sigma_x^2 + (1 + \alpha\gamma)\sigma_{x,s}}{\sigma_s^2 + \alpha^2\sigma_x^2 - 2\alpha\sigma_{x,s}}.$$

Here σ_x^2 and σ_s^2 are the variances of x and s , and $\sigma_{x,s}$ is their covariance. The conditions under which the regression coefficient will reveal γ are: (1) $\alpha = 0$ (no contamination of unemployment as a measure of slack) and (2) $\sigma_{x,s} = 0$ (structural shifts uncorrelated with tightness). These conditions

are familiar from standard simultaneous-equations econometrics—ordinary least squares estimation is okay in a simultaneous system if it is triangular (one-way causation) and the disturbances are uncorrelated.

There is no compelling reason to suspect correlation of slack and structural shifts, but the implicit assumption that $\alpha = 0$ is fundamentally questionable. Many observers have doubted the reliability of unemployment as a measure of slack and explicitly raised the concern that unemployment fails to measure slack correctly when structural shifts in participation occur. Research based on regressing participation on unemployment cannot answer the important question of hidden unemployment, and the presence of hidden unemployment (in the sense of nonzero values of α) biases estimates of the effect of slack on participation.

The paper thus leads with its strength, direct information about sources of structural changes in participation. This information is not based on regressions of participation on unemployment. This part of the paper provides evidence on reported reasons for nonparticipation, direct measurement of demographic indexes, and direct measurement of education enrollment, disability and disability benefits, and retirement.

With respect to the often-noted role of the changing age composition of the working-age population, which by itself would account for an important part of the decline in participation, the paper notes, though only in a footnote, Robert Shimer's finding that other trends in composition almost exactly cancel the downward trend in participation from the age composition. In my opinion, Shimer's approach, based on considering all composition effects together, gives a clearer picture. That picture features a larger role for special forces operating on participation, because it finds no net effect from the various compositional sources of change in participation. The model presented toward the end of the authors' paper does make composition adjustments across multiple characteristics of the population through cohort effects.

The paper focuses almost exclusively on the standard unemployment rate as a measure of labor-market slack. Although I believe that current unemployment, at 5.6 percent, is a reasonable basis, by historical standards, for concluding that the slack from the crisis is now gone, I think it is important to recognize the heterogeneity across a range of measures of slackness.

On the one hand, short-term unemployment—the fraction of the labor force who became unemployed within the past 6 weeks—is currently remarkably low. At 1.6 percent, it is lower than ever before recorded. This measure of unemployment was 1.7 percent in the strong labor market

of 2007, just before the crisis, when the overall unemployment rate was a robust 4.6 percent, and it was 1.8 percent in the even stronger labor market of 2000, when the unemployment rate was 4.0 percent. Another measure showing an exceptionally strong market is the average time taken by employers to fill jobs. Longer recruiting times indicate that the condition of the labor market is favorable to job seekers and correspondingly more difficult for employers to match with those job seekers. At 28 days, average duration is the same as in the strong market of 2007 and longer than the 26 days recorded in 2001, a year of low (4.8 percent) unemployment.

On the other hand, long-term unemployment, a legacy of the wave of deep job loss from the crisis, remains above normal. In 2014, workers still searching after 6 months of unemployment accounted for 2.1 percent of the labor force, down from a peak of double that level in 2010 but above the normal level of about 1 percent of the labor force. Another indicator showing remaining slack in the labor market is the fraction of workers who would choose full-time work if it were available, but are now working part time. At 3.0 percent, it is above its normal level of about 2 percent.

The authors follow an earlier team from the Federal Reserve—Christopher Erceg and Andrew Levin, cited in the authors' paper—in studying the relation between participation and unemployment in a state-level panel. Although the findings in the new paper suggest a weaker negative correlation between the two variables than Erceg and Levin found, the correlation is quite a bit stronger than in national data. There is an emerging puzzle in labor-macro research between studies at the state level and those at the national level. The most extreme example is a widely discussed paper by Marcus Hagedorn, Iourii Manovskii, Fatih Karahan, and Kurt Mitman (2015), which finds huge effects of unemployment insurance on unemployment at the state level. That paper uses a regression-discontinuity design to try to overcome simultaneity issues, yet finds a relationship that does not appear in national data.

The paper under discussion here concludes with an ambitious study of participation rates at a high level of disaggregation, using data from the Current Population Survey (CPS). The authors' idea is to exploit these rich data with an unrestrictive specification, with a full set of cohort effects and age effects, and with separate models for men and women. They chose to model time effects as a linear combination of observed variables. An alternative would have been to estimate a full set of time effects—that is, substituting time dummies for the observed variables. A full set is identified, except for a linear trend. The authors impose a strong condition, not needed for identification, by saying that the time effects have to track their

set of time-varying variables. But the only restriction on the time effects required for identification is to specify the trend, say to zero.

The authors were kind enough to calculate a set of results for me that dropped the time-varying variables and replaced them with a complete set of year dummies. For women, the restriction they impose looks fairly good—the difference between the constrained and unconstrained time effects is not too big. The time-varying variables, including unemployment, are capable of mimicking the actual underlying time effects. However, for men the fitted values from the authors' constrained specification are quite different from the freely estimated time effects, a sign that there is a problem with the specification. The authors plan to pursue this issue. They note that the cohort effects for men are also large in unexpected ways.

My earlier remarks on the authors' reliance on unemployment as the primary measure of conditions in the labor market also apply to the results from the CPS. Among the time-varying variables used to capture the time effects, only the unemployment gap—actual unemployment less the Congressional Budget Office's measure of the natural unemployment rate—comes from the labor market. The CPS model is not equipped to deal with the question of whether the unemployment rate today gives too favorable a view that no slack remains in the labor market.

In summary, this paper advances research on labor-force participation in a number of useful ways. But much more remains to be done. The contrast between the results from national and state-level data on the correlation between unemployment and participation is a major challenge. In all approaches to studying the causal relations between the two variables, we need to go beyond correlations and regressions to try to identify the structural parameters.

Monetary policymakers are paying close attention to conditions in the labor market. Under its dual mandate, the Federal Reserve needs to determine what constitutes full employment, and the current state of the labor market seems to be close to that point. This paper concludes that the major decline in participation should not be interpreted to indicate that slack is greater than the unemployment rate signals. Notwithstanding my misgivings about some of the econometrics in the paper, I agree with this conclusion. Does this mean that the time has come for contractionary monetary policy? I think not. The Federal Reserve became concerned about the labor market overheating in 1994, but ultimately let the unemployment rate reach 3.8 percent. This seemed extraordinarily tight, but there was no explosion of inflation. The unemployment rate hit 4.4 percent in 2007, again with no

inflation issue. Today, inflation forecasts are at levels below target, so there seems no reason to resist further declines in unemployment.

A final note: The current discussion of monetary policy has completely ignored the compelling arguments of Michael Woodford (2013), Gaudi Eggertsson (2003), and others that policy should deliberately generate *extra* inflation and unusual tightness right after the economy escapes the zero lower bound, because doing so would help prevent harm from the bound in future episodes.

REFERENCES FOR THE HALL COMMENT

- Eggertsson, Gaudi, 2003. "How to Fight Inflation in a Liquidity Trap: Committing to Being Irresponsible." Working Paper no. 03/64. Washington: International Monetary Fund.
- Erceg, Christopher J., and Andrew T. Levin. 2013. "Labor Force Participation and Monetary Policy in the Wake of the Great Recession." Working Paper no. 13/245. Washington: International Monetary Fund.
- Hagedorn, Marcus, Iouri Manovskii, Fatih Karahan, and Kurt Mitman. 2015. "Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects." Working Paper no. 19499. Cambridge, Mass.: National Bureau of Economic Research.
- Woodford, Michael. 2013. "Monetary Policy Targets after the Crisis." Paper presented at the International Monetary Fund conference, "Rethinking Macro Policy II." Washington, April 16.

COMMENT BY

JAMES H. STOCK¹ It is a pleasure to discuss this important paper on the timely and pressing topic of the reasons for the decline in the labor force participation rate (LFPR). Members of this team of authors have been on the forefront of using cohort data to examine trends in the LFPR, and their previous methodological and substantive contributions, notably in Stephanie Aaronson and others (2006), have been at the center of the burgeoning literature on the reasons for the recent sharp decline in the LFPR. The paper they present in this volume updates and goes beyond their previous work in a number of ways, including additional modeling of the structural trends underlying the decline in the LFPR, so as better to forecast the participation rate's future path.

As the macroeconomics community is now realizing, the evolution of the LFPR is central to the economy's growth prospects over the medium

1. I am grateful to John Coglianese for his assistance in preparing these comments.

and longer run. That is because the growth rate of the LFPR enters one-for-one in the growth rate of GDP: at any given level of the unemployment rate, a larger labor force means more workers, which means more output. Because the LFPR is declining, so is the underlying supply-side growth rate of GDP. This slowdown in overall rate of economic growth has important implications. First, it poses substantial challenges to monetary policy. In particular, if the slower growth of GDP is accompanied by a lower equilibrium real interest rate (as is consistent with historical evidence), then the equilibrium nominal rate is closer to zero. Consequently, absent a compensating upward adjustment to the inflation target, it becomes more likely that monetary policy will be constrained by the zero lower bound in future recessions. With hamstrung monetary policy, future recessions could be deeper, and recoveries longer, than they otherwise would be with higher trend GDP growth or, alternatively, with a higher inflation target.

A second implication of the slowdown in GDP growth is that the trend decline in the LFPR is a key reason why this recovery has been slow. The Council of Economic Advisers (CEA) (2013) looked at a number of studies and estimated that one-half to two-thirds of the current slow recovery, relative to the 1980s, was attributable not to shocks associated with this recession but to slower long-term trend GDP growth (also see Stock and Watson 2012, CBO 2012, and Hall 2014). A third implication is that slower GDP growth implies slower growth in tax receipts, all else equal, which if unmatched by a reduction in the growth rate of spending will impose additional fiscal burdens and additionally constrain future countercyclical fiscal policy (see Gordon 2014).

With these issues as a backdrop, the present paper by Aaronson and others tackles two questions: why the LFPR has been declining and whether that decline can be expected to continue. The authors present a decomposition of the decline into structural trend factors, cyclical factors, and other factors (residual). Their approach combines several long-term factors into the trend. While there is a high degree of confidence in at least one of these factors—the retirements of aging baby boomers—other aspects of these trends are less well identified or are unidentified and simply enter the cohort effects. In addition, I will argue that the cyclical components of their model are also not cleanly separated from the trend components.

In these comments, I therefore present an alternative decomposition into a well-identified pure aging trend, a historical cyclical component (specifically, a cyclical component based on pre-2007 cyclical patterns in the LFPR), and a residual that contains a number of possible factors, some of which might be persistent and some of which might be special features of

the deep recession and long recovery. This alternative decomposition follows and extends CEA (2014) by estimating cyclical components for demographic subgroups, providing standard errors for cyclical components, and updating the CEA decomposition using current data. Using this alternative decomposition framework, I reach four main conclusions.

First, of the 3.2-percentage-point decline in the LFPR since 2007Q4, roughly half (1.6 points) is due to a pure aging effect.² As I discuss below, this pure aging component is well identified.

Second, the historical cyclical component, which is also well identified, was important during the recession and the first few years of the recovery. Now, however, the cyclical contribution is substantially diminished, and as of 2014Q2 it accounted for approximately -0.5 percentage point, or one-sixth of the decline; a 95-percent confidence interval for this cyclical component is $(-0.3, -0.7)$ percentage point. Using data through 2014Q4, the historical cyclical contribution is less, -0.4 percentage point.

Third, young adults have had a particularly large decline in the LFPR. While a large portion of this decline was initially accounted for by historical cyclical factors, only 0.6 (standard error 0.4) of the 3.8-percentage-point decline in participation among 20–24-year-olds from 2007Q4 to 2014Q2 is accounted for by cyclical factors.

Fourth, the remaining residual, 1.1 percentage points in 2014Q2, comprises approximately one-third of the decline in the overall LFPR. Time series correlations suggest that some, possibly much, of this residual could disappear as the labor market normalizes. However, both time series and cohort analyses are limited in what they can say about the structural sources of this residual. In the final section of these comments I suggest that different methods are needed to understand and to project this residual.

In the remainder of my discussion I will give a brief summary of the aging trend, a longer discussion of the cyclical component (including standard errors), including a comparison of the cyclical components for different age/gender subgroups, and a discussion of the residual.

PURE AGING TREND. The aging trend arises because of the changing age distribution of the population, in particular the entrance of the baby boom into retirement years. If the cohort effects were well identified—that is, if they were the composite of individual fixed effects that reflected

2. The authors use data adjusted to account for changes in population controls, and with their data the decline is 2.8 percentage points (see their table 1), while I use the published unadjusted data; remaining discrepancies are due to rounding.

immutable additive individual attitudes toward work—then one would want to allow for cohort-induced changes to the age-participation profile. However, linear cohort effects are not identified separately from linear time effects, a point that is particularly important in the current context because of the essentially linear trend decline in the male LFPR since the 1970s (see the authors' figure 4).


The authors address this identification problem by assuming that the included time-varying variables in their econometric model capture all the time effects. This normalization is not testable, however, so any misspecification in linear time effects (trends) is loaded into the cohort effects. Indeed, this unavoidable conflation of time trends and trends in the cohort effects may be one reason why the earlier version of this model (Aaronson and others 2006) predicted that the out-of-sample decline in the LFPR, specifically that model's evolving cohort effects, could have captured and extrapolated the in-sample trend. Thus, while their method is appealing, allowing for cohort effects in estimating the aging trend opens a host of issues.

For this reason, the aging trend that I report here freezes the participation rates at their 2007Q4 values and allows the age distribution to change. This "pure aging" trend, which is shown for the total LFPR in my figure 1, accounts for 1.6 percentage point of the 3.2-percentage-point decline in the LFPR.³

CYCLICAL COMPONENT. The complication in estimating the historical cyclical component is that it must be done in the presence of complicated slowly varying trends. Moreover, those trends must somehow be distinguished from the cyclical component during the part of the data of greatest interest, the decline during the recession and the long recovery. This complication is not fully resolved by using state panel data, as are used both in the authors' paper and in other papers in the literature. To the extent that there are different nonlinear state trends, using time fixed effects and state linear trends (as in the panel data regressions in the authors' table 2) leaves the trend component misspecified.

3. The authors estimate the pure aging trend to account for 1.3 percentage points of the decline. The discrepancy between their estimate and my 1.6 percentage points estimate here is likely due to their use of chain weighting. In this application, my approach has the virtue of using pre-recession participation rates, whereas chain weighting uses participation rates that are influenced by the recession; thus the approach used here answers the question of what the aging effect on the LFPR would have been, absent the recession.

Figure 1. Labor Force Participation Rate (LFPR), Men and Women, Ages 16+, 2007Q4–2014Q4



Source: Bureau of Labor Statistics and author’s calculations.

a. The pure aging trend is the participation rate predicted by the changing age distribution, fixing age-specific participation rates at their 2007 level. Trend and trend + cycle are relative to 2007Q4.

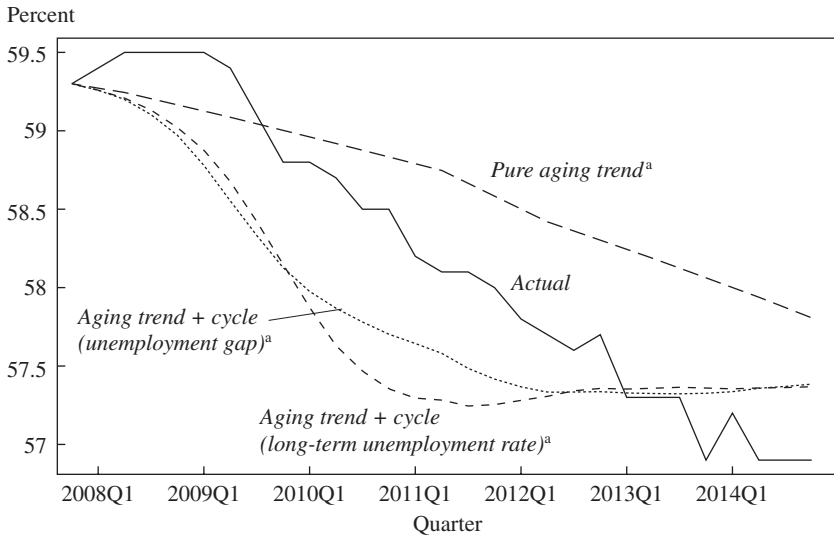
b. The cyclical component is estimated using 1970–2007 data, a flexible trend, and a distributed lag of the unemployment slack measure (either the CBO unemployment gap or the long-term unemployment rate).

This complication can be resolved, however, by using time series data and simultaneously estimating a flexible nonparametric trend and a cyclical component specified in terms of a distributed lag of a measure of cyclical slack, as well as by estimating the historical cyclical component using coefficients estimated using data through 2007Q4. To this end, I consider an extension of the approach used by the Council of Economic Advisers (2014), in which the growth rate of the LFPR is expressed as a flexible trend μ_t , a distributed lag of CBO’s unemployment gap u_t^{gap} , and a remainder v_t , which in general is serially correlated:

$$\Delta \ln LFPR_t = \mu_t + \beta(L) \Delta u_t^{gap} + v_t.$$

This is the partially linear regression model of Robinson (1988) and Stock (1989). The trend term μ_t needs additional structure to be identified; here, this is done by treating it as the estimand of a long time-series filter. The

Figure 2. LFPR, Women, Ages 16+, 2007Q4–2014Q4




Source: See the notes to figure 1.
 a. Trend and trend + cycle are relative to 2007Q4.

specific approach used here estimates the partially linear regression model using a biweight filter.⁴

The lag polynomial $\beta(L)$ was estimated using data through 2007Q4 so that the estimated cyclical component represents the historical cyclical component. To the extent that the Great Recession induced additional cyclical movements in the LFPR above and beyond those captured by the pre-2007Q4 data, they would appear in the residual \hat{v}_t .

4. Let y_t and x_t respectively denote $\Delta \ln LFPR_t$ and u_t^{gap} . The algorithm proceeds as follows: (i) estimate the univariate trends in y_t and x_t using a symmetric kernel (filter); (ii) deviate y_t and x_t from those trends, and call these respectively \tilde{y}_t and \tilde{x}_t ; (iii) regress \tilde{y}_t on the chosen distributed lag of \tilde{x}_t to obtain the estimator $\hat{\beta}(L)$; (iv) compute the cyclical component $\hat{\beta}(L)x_t$ and the cyclically adjusted y_t as $y_t^{CA} = y_t - \hat{\beta}(L)x_t$; (v) obtain the estimator of the trend $\hat{\mu}_t$ as the smoothed value of y_t^{CA} , smoothed using a symmetric kernel. Step (v) does not affect the estimate of the cyclical component. For step (ii), a biweight filter with a window width of 40 quarters was used. For step (iii), several lag specifications were examined; the results here are for the four-quarter change in the unemployment gap and its 4- and 8-quarter lags (a 12-quarter unrestricted distributed lag yields similar, but noisier, estimates of the cyclical component). Standard errors for the cyclical component at date t are computed as the standard error of the linear combination of the $\hat{\beta}$ s indicated by the relevant values of x_t , using Newey-West standard errors (12-quarter truncation parameter). For recent work on this model see Cai (2007) and Zhang and Wu (2012).

Figure 3. LFPR, Men, Ages 16+, 2007Q4–2014Q4



Source: See the notes to figure 1.
 a. Trend and trend + cycle are relative to 2007Q4.

My figures 1–3 present the LFPR since the cyclical peak in 2007Q4, the aging trend, computed as described above, and the sum of the aging trend and the cyclical component, both for the total LFPR and separately for women and men. The figures present two estimates of the historical cyclical component, one that uses the CBO unemployment gap as the cyclical measure (as in the equation above) and an alternative that uses the long-term (≥ 27 weeks) unemployment rate. In both cases the coefficients are estimated using quarterly data from 1970Q1 through 2007Q4. For both men and women, using the long-term unemployment rate as the cyclical measure results in a larger estimated cyclical component in 2010 and 2011, but by 2014 the two measures estimate very similar cyclical components. For brevity, the discussion below focuses solely on the historical cycle estimated using the more conventional of these two measures, the CBO unemployment gap.

For women, the decline through 2013 is more than explained by the aging trend plus the historical cycle, and through 2014Q2 all but 0.5 percentage point of the 2.4-percentage-point decline is explained by these two components. In contrast, for men the trend and cyclical components account for

Table 1. Change in LFPR Due to Cyclical Component, 2007Q4–2010Q4

Age group	<i>Men and women</i>		<i>Men</i>		<i>Women</i>	
	<i>Total</i>	<i>Cycle^a</i>	<i>Total</i>	<i>Cycle^a</i>	<i>Total</i>	<i>Cycle^a</i>
16+	-1.6	-1.1 (0.2)	-2.2	-0.9 (0.2)	-0.8	-1.1 (0.2)
16–19	-6.6	-3.5 (0.7)	-5.6	-3.6 (0.8)	-7.6	-3.4 (0.8)
20–24	-3.0	-1.6 (0.4)	-4.6	-1.8 (0.6)	-1.3	-1.7 (0.6)
25–54	-1.0	-1.0 (0.2)	-1.8	-0.5 (0.2)	-0.3	-1.4 (0.3)
55+	1.2	-0.1 (0.3)	0.9	-0.6 (0.4)	1.5	0.3 (0.4)

Source: Bureau of Labor Statistics and author's calculations.

a. The cyclical component is estimated using data from 1970–2007 using a constrained three-year lag of the unemployment gap and a flexible trend, where the trend is estimated using a long moving average (a biweight filter as described in the text). Standard errors for the cyclical component appear in parentheses.

the early drop in the LFPR (through 2010), but as of 2014Q2 the aging trend accounts for 1.8 percentage points, or half, of the 3.9-percentage-point decline since 2007Q4, and the cyclical component accounts for only 0.4 percentage point of the decline, leaving more than 40 percent of the decline, or 1.7 percentage points, as the unexplained residual. Combined, the last two years of decline in the overall LFPR are not explained by the aging trend and historical cyclical component. Updated through 2014Q4, for the overall LFPR the historical cyclical component accounts for only 0.4 percentage point of the decline, leaving 1.1 percentage points, or one-third of the decline, as the residual.

My table 1 presents the resulting estimated cyclical components as of 2010Q4 for the total LFPR and for various age and gender groups. My table 2 presents the comparable results as of 2014Q2. As of 2010Q4, all the decline in the LFPR for workers ages 25–54 was cyclical, but by 2014Q2 only one-fourth of the decline for workers ages 25–54 was cyclical.

RESIDUAL. I have argued that the aging trend and the historical cyclical components are well identified and well understood. The aging trend will remain with us as the population ages, and the cyclical component is rapidly vanishing as the labor market approaches full recovery. What this decomposition highlights is that there is a large part of the decline—approximately 1.1 percentage points for the overall LFPR—which is accounted for neither by pure aging nor by the waning cyclical component.

Table 2. Change in LFPR Due to Cyclical Component, 2007Q4–2014Q2

Age group	<i>Men and women</i>		<i>Men</i>		<i>Women</i>	
	<i>Total</i>	<i>Cycle^a</i>	<i>Total</i>	<i>Cycle^a</i>	<i>Total</i>	<i>Cycle^a</i>
16+	-3.2	-0.5 (0.1)	-3.9	-0.4 (0.1)	-2.4	-0.6 (0.2)
16–19	-7.6	-0.7 (0.5)	-7.8	-0.6 (0.6)	-7.5	-0.9 (0.7)
20–24	-3.8	-0.6 (0.4)	-5.1	-0.8 (0.4)	-2.4	-0.5 (0.5)
25–54	-2.1	-0.6 (0.2)	-2.5	-0.3 (0.1)	-1.5	-0.8 (0.2)
55+	1.1	-0.2 (0.2)	0.6	-0.4 (0.2)	1.4	0.0 (0.2)

Source: Bureau of Labor Statistics and author’s calculations.

a. The cyclical component is estimated with data from 1970–2007 using a constrained three-year lag of the unemployment gap and a flexible trend, where the trend is estimated using a long moving average (a biweight filter as described in the text). Standard errors for the cyclical component appear in parentheses.

As shown in my table 1 (which controls, crudely, for shifting demographics by estimating the decline within broad age groups), the non-cyclical components of the decline vary substantially across demographic subgroups. In particular, according to these estimates, the large declines in the LFPR for workers ages 20–24 initially had large cyclical components, but they account for less than one-fifth of the decline through 2014Q2. A particularly striking feature of table 1 is that, for young men ages 20 to 24, only 0.8 point of the large 5.1-percentage-point decline in participation through 2014Q2 is attributed to historical cyclical patterns. Looking forward, a key question is whether being a full-time student, a status that has risen sharply over this episode, is an accommodation to the lack of jobs or part of an ongoing trend of increasing education among younger potential workers.

The Council of Economic Advisers (2014) uses time series methods to suggest both that some of the residual is associated with the severe impacts of this recession on the labor market, and is likely to disappear, and that Social Security Disability Insurance plays at most a small role in explaining the residual. However, omitted variables are inevitable, so causal conclusions based on time series data are unlikely to be compelling in this case. While the cohort model in the paper under discussion might initially seem more promising, it too faces the fundamental identification problem of distinguishing cohort effects from time effects. Indeed, in the decomposition

in their table 3, cohort effects “explain” 1.7 percentage points—more than half—of the decline in the overall LFPR from 2007Q4 to 2014Q2; however by construction these cohort effects are simply things otherwise unaccounted for in the model. In addition, some of the proximate determinants of participation in their model are actually endogenous with decisions that are likely affected by the same factors that determine participation, for example the years of schooling (and the decision whether to attend school full-time) of younger adults.

For these reasons, it seems more promising to use individual-level data with state and individual-level variation to identify the factors influencing labor force decisions. While research using individual-level data is the norm in the labor economics community, there remains much room to integrate that research, which tends to focus on specific channels, with the broader perspective in the LFPR literature. To give just two recent examples that focus on the timing of retirement and use individual-level data: Alan Gustman and Thomas Steinmeier (2014) estimate substantial impacts of health on the timing of retirement, and Robin Lumsdaine and Stephanie Vermeer (2014) suggest that the arrival of a new grandchild increases the probability of a working woman retiring. This latter estimate suggests that the trend toward later childbearing could be one factor behind the increasing retirement age of older women, and it could potentially revert to the extent that childbearing was postponed during the recession and slow recovery.

In conclusion, the authors have done a significant service by extending their earlier work and integrating it with state-level analysis to provide a range of estimates of the aging trend and cyclical components of the decline in the LFPR. Their estimates are in the range of most of the others in the literature and in the range of those I provide here. The methods I use here have the advantage of identifying the historical cyclical effect from pre-2007 data in a way that allows for, and simultaneously estimates, a flexible cyclically adjusted trend.

Looking across the results in the authors’ paper and the alternative estimates I present here, the picture of the participation rate is one of continuing decline over the coming decade, possibly with some near-term sideways motion as the economy continues to strengthen. Some results point to the possibility that the remaining large residual decline in the LFPR, approximately one percentage point, will be a transitory consequence of the vast disruptions imposed on the lives of workers by the Great Recession. But the sources of this additional decline remain only partially understood, and future research is needed to sharpen projections.

REFERENCES FOR THE STOCK COMMENT

- Aaronson, Stephanie, Bruce Fallick, Andrew Figura, Jonathan Pingle, and William L. Wascher. 2006. "The Recent Decline in Labor Force Participation and Its Implications for Potential Labor Supply." *Brookings Papers on Economic Activity* 1.
- Cai, Zongwu. 2007. "Trending Time-Varying Coefficient Time Series Models with Serially Correlated Errors." *Journal of Econometrics* 136: 163–88.
- Congressional Budget Office (CBO). 2012. "What Accounts for the Slow Growth of the Economy after the Recession?" Washington.
- Council of Economic Advisers. 2013. "Economic Report of the President: Transmitted to the Congress March 2013, Together with the Annual Report of the Council of Economic Advisers." Washington.
- . 2014. "The Labor Force Participation Rate since 2007: Causes and Policy Implications." Washington: Executive Office of the President.
- Gordon, Robert J. 2014. "A New Method of Estimating Potential Real GDP Growth: Implications for the Labor Market and the Debt/GDP Ratio." Working Paper no. 20423. Cambridge, Mass.: National Bureau of Economic Research.
- Gustman, Alan L., and Thomas L. Steinmeier. 2014. "The Role of Health in Retirement." Working Paper no. 19902. Cambridge, Mass.: National Bureau of Economic Research.
- Hall, Robert E. 2014. "Quantifying the Lasting Harm to the U.S. Economy from the Financial Crisis." *NBER Macroeconomics Annual 2014*, forthcoming.
- Lumsdaine, Robin L., and Stephanie J. C. Vermeer. 2014. "Retirement Timing of Women and the Role of Care Responsibilities for Grandchildren." Working Paper no. 20756. Cambridge, Mass.: National Bureau of Economic Research.
- Robinson, Peter M. 1988. "Root-N-Consistent Semiparametric Regression." *Econometrica* 56, 931–54.
- Stock, James H. 1989. "Nonparametric Policy Analysis." *Journal of the American Statistical Association* 84, no. 406: 567–75.
- Stock, James H., and Mark W. Watson. 2012. "Disentangling the Channels of the 2007–2009 Recession." *Brookings Papers on Economic Activity*, Spring: 81–135.
- Zhang, Ting, and Wei Biao Wu. 2012. "Inference of Time-Varying Regression Models." *Annals of Statistics* 40, no. 3: 1376–1402.

GENERAL DISCUSSION Robert Gordon began the discussion by drawing an analogy between this paper and Louise Sheiner's conference paper on the geography of medical expenditures,¹ both of which he thought would soon become classics in the literature. The variation in health expenditures,

1. "Why the Geographic Variation in Health Care Spending Cannot Tell Us Much About the Efficiency or Quality of Our Health Care System" (published in this volume of *Brookings Papers*; see page 1.)

he thought, was due to the concentration of low-income African-Americans living in the Cotton Belt and in Chicago since the Great Migration. Gordon thought that the same factors were at work in the labor force participation rate.

Agreeing with discussant James Stock's assessment that the decline in the labor force participation rate was structural, Gordon shared some findings on potential GDP that he had recently published as an NBER working paper. He had found that between 1972 and 2004, the GDP growth rate had averaged 3.2 percent per year, but in the last decade it had averaged only 1.6 percent. Fully 0.9 percentage point of the slowdown could be attributed to the decline in the labor force participation rate. Gordon believed this decline in potential GDP had significant implications. Noting that the Congressional Budget Office (CBO) has already lowered its expectation of 2024 potential GDP by \$2 trillion, he said his own estimates suggested it should be lowered by an additional \$2 trillion.

Justin Wolfers thought that terms like "structural" and "cyclical" were loosely defined in the literature. The authors of this paper used the term *cyclical* to refer to factors affecting the labor force participation rate that will tend to be responsive to monetary policy. But that definition is different from one that is affected by changes in disability rules. Wolfers thought the authors got the definition right, but worried that a lot of other people have been getting it wrong. Admiring the elegance of the authors' charts, he said that the model used by the authors fits the data almost perfectly both because of the large number of parameters and also because both age and cohort effects are included. The nice fit says a lot about the mathematics of the models but, in Wolfers' view, it tells one nothing about the actual labor market.

Wolfers agreed with Gordon that the prospect for a recovery in the labor force participation rate depends both on the estimated cyclicity of the rate and on where the economy is in the cycle. The unemployment rate is currently 6.1 percent, and the CBO estimated the natural rate of unemployment to be 5.5 percent; if the unemployment rate were to fall further, so that the gap between the unemployment rate and the natural rate became zero, there would be no prospect for a recovery in the labor force participation rate, simply because there would be no cycle remaining, not because the rate itself was noncyclical.

Kristin Forbes agreed with discussant Robert Hall that the unemployment gap would be a better measure to study than the unemployment rate. Adding to the discussion about how the terms "structural" and "cyclical"

should be defined, she opined that the line between what could be affected by the monetary policy and what could not was blurrier than implied in the paper. As an example, she pointed out that the aging of the labor force would typically be thought of as structural, but in the United Kingdom a number of elderly Britons were now returning to the labor force. Changing trends of this kind made it difficult to define the difference between structural and cyclical factors. Referring to an analysis done by Goldman Sachs, she mentioned that many of the concerns expressed by the discussants—cohort effects, lag structure, and time effects among them—had been accounted for in a similar model, which found that these modifications could significantly change the results on the relative importance of cyclical and structural factors. She asked if the authors had seen the Goldman analysis and what they thought of it.

Katharine Abraham noted that people who are out of the labor force for what appear to be structural reasons can be drawn back in if job opportunities are sufficiently good, reinforcing the point that it may be difficult to distinguish empirically between structural and cyclical factors.

Jonathan Pingle suggested that the cohort effects were essentially unobserved components in the model. He had used similar models to avoid endpoint bias, a problem for many time-series filters. Pingle had found that cohort effects were well identified, even though other trends may have been picked up by the effects, because differential life-cycle labor force attachment was embedded in the cohort effects. For the recent downturn to unduly influence a cohort effect, it needed to be identified primarily by the recent business cycle. Thus, only the newer birth cohorts were at risk. In Pingle's view, the authors' assumptions did a good job of dealing with that, and the paper showed the effects of alternative assumptions.

Valerie Ramey suggested that comparing the labor force participation rate across countries would be informative. She noted that the participation rate among the 15-to-64 age group had risen by 1.7 percentage points in the United Kingdom while simultaneously declining by 1.3 percentage points in the United States. She was reminded of an explanation Steven Davis raised at a previous meeting of the Brookings Panel, to the effect that the large number of persons with criminal records in the United States depressed labor force participation.

Steven Davis added some numbers to flesh out Ramey's comment on Americans with a criminal history. He said the cumulative risk of having spent some time in prison by age 30 to 34 for white men with a high school

education had risen from 4 percent in 1979 to 28 percent in 2009. For black men, both the increase and the level were much larger. Davis found the paper's implications for future economic growth to be sobering. Aging alone is likely to reduce the labor force by 1 to 1.5 percentage points over the next decade. The secular decline in labor market fluidity that he and John Haltiwanger documented in a paper presented to the Jackson Hole conference is a force behind declining employment within age groups.

Haltiwanger spoke up to add that declines in labor market fluidity are not only related to declines in the labor force participation rate but also likely related to declines in real wage and productivity growth. The decline in fluidity makes the labor market a thinner market with fewer opportunities to make matches.

Andrew Levin mentioned a study by the Pew Foundation which found that the number of mothers staying at home because they could not find work had increased dramatically, from 1 percent to 6 percent. He thought it would be useful for the Federal Reserve to survey prime-age nonworking individuals. Referencing the large increase in part-time employment, which has recovered about half of the decline experienced following the recession, Levin thought that this was evidence of a shortage in aggregate demand.

Betsey Stevenson suggested that the decline in participation for stay-at-home mothers might be related to the point just made by Davis and Haltiwanger. As the fluidity of the market declines, the offer rate declines for these women. Linking this story to the international comparisons mentioned during the discussion, she noted that the United States had gone from being a leader in female labor force participation to being a laggard.

Speaking on behalf of the paper's six coauthors, Christopher Smith joined with the discussants who were concerned about the unemployment rate being a distorted signal of slack in the labor market. He noted that identifying differences in cyclicalities based on a single event was difficult by definition. But he appreciated many of the suggestions to address cyclicalities by not confining it to one episode. He added that measuring the uptake of food stamps and the duration of unemployment might be useful extensions to make.

Referring to the Goldman Sachs model that Forbes had mentioned, which looked at panel regressions, Smith said the authors' model, which modeled cohort effects, was quite similar apart from the specifications that were emphasized. He concluded that it was a simple question of choosing which specifications were to be preferred, adding that there was a lot of uncertainty in the choosing.

Smith agreed that the cohort effects provided only limited information—like the residual to any regression. He said he and the coauthors would continue to try and explain the cohort effect by looking at fluidity, prison, polarization, and other factors. International comparisons could illustrate some of these other factors, although they would be difficult to execute due to the challenges of varying definitions, survey methods, and measurements among the different countries. The authors had not yet attempted to do such a comparison, but Smith thought it would be possible.

Regarding the terminology of “cyclical” and “structural,” he acknowledged that keeping clear definitions was indeed difficult, not only for the reasons stated by the discussants but also because the cyclical factors could become permanent if they lasted long enough. When that happened, Smith asked, what should those factors then be called?