CHRISTOPHER J. ERCEG

ANDREW T. LEVIN

Labor Force Participation and Monetary Policy in the Wake of the Great Recession

This paper provides compelling evidence that cyclical factors account for the bulk of the post-2007 decline in the U.S. labor force participation rate (LFPR). We then formulate a stylized New Keynesian model in which the LFPR is practically acyclical during "normal times" but drops markedly following a large and persistent aggregate demand shock. These considerations have potentially crucial implications for the design of monetary policy, especially when interest rate adjustments are constrained by the zero lower bound; specifically, monetary policy can induce a more rapid recovery of the LFPR by allowing the unemployment rate to fall below its natural rate.

JEL codes: E24, E32, E52, J21
Keywords: New Keynesian models, unemployment rate, simple monetary policy rules, zero lower bound.

A LONGSTANDING AND WELL-ESTABLISHED fact in labor economics is that aggregate labor supply has been only mildly procyclical over the past half-century (cf. Mincer 1966, Pencavel 1986, Heckman and Killingsworth 1986). Consequently, macroeconomists have largely focused on the unemployment rate
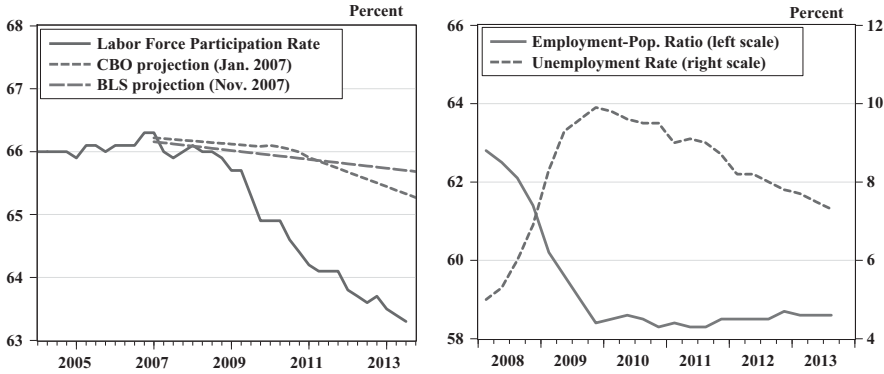
FIG. 1. Labor Market Developments in the Wake of the Great Recession.

NOTES: In the left panel, the solid line denotes the actual evolution of the LFPR from 2004:Q1 to 2013:Q2; the short-dashed line denotes the CBO's January 2007 projection for the potential LFPR, that is, the path of LFPR that would be consistent with full employment; and the long-dashed line denotes the LFPR projection that was published by the BLS in November 2007. In the right panel, the solid line denotes the ratio of employment to the civilian noninstitutionalized population aged 16 years and older (left scale), and the dashed line denotes the civilian unemployment rate (right scale).

as a business cycle indicator while abstracting from movements in labor force participation.[1] Similarly, the literature on optimal monetary policy and simple rules has typically assumed that unemployment gaps and output gaps can be viewed as roughly equivalent (Orphanides 2002, Taylor and Williams 2010).

In this paper, we reconsider such conventional wisdom in light of labor market developments since the Great Recession. As shown in the left panel of Figure 1, the labor force participation rate (LFPR) has fallen about 2.5 percentage points—a striking contrast to the relatively modest declines that were projected by the Congressional Budget Office (CBO) in January 2007 and by the Bureau of Labor Statistics (BLS) in November 2007, just prior to the onset of the recession.[2] Moreover, as shown in the right panel, the employment-to-population ratio has remained close to its postcrisis trough of about 58.5%. In effect, the notable decline in the unemployment rate— which has retraced more than halfway back from its peak—mainly reflects a high incidence of departures from the labor force.

Our paper provides compelling empirical evidence that cyclical factors account for the bulk of the recent decline in the LFPR. We then proceed to formulate a stylized New Keynesian model in which households' labor market exit and reentry decisions are associated with significant adjustment costs. Our model analysis highlights how policy rules that respond to broader measures of labor market slack that include the

---

1. For example, the now-classic paper by Nelson and Plosser (1982) analyzed the time series behavior of an array of macroeconomic indicators, including aggregate employment and the unemployment rate, but did not consider any measure of labor force participation.

2. The labor force projections published by the Social Security Administration had broadly similar contours. In contrast, as of January 2007, the Federal Reserve Board's (FRB) internal staff forecast for 2007–08 embedded a fairly steep decline in the trend participation rate, which is consistent with the empirical analysis of Aaronson, Park, and Sullivan (2006); sensitive FOMC materials remain confidential for a 5-year period, so the FRB staff forecasts prepared in 2007 were released to the public in early 2013.

cyclical component of participation may have very different implications for how the economy recovers from a deep recession than "standard" rules that focus on the unemployment gap.

More specifically, our analysis of state-level employment data indicates that cyclical factors can fully account for the post-2007 decline of 1.5 percentage points in the LFPR for prime-age adults (i.e., 25–54 years old). We define the labor force participation gap as the deviation of the LFPR from its potential path implied by demographic and structural considerations, and we find that as of mid-2013 this gap stood at around 2%. Indeed, our analysis suggests that the labor force gap and the unemployment gap each accounts for roughly half of the current employment gap, that is, the shortfall of the employment-to-population rate from its precrisis trend.[3]

Our empirical analysis is broadly consistent with a number of other recent studies. Aaronson, Davis, and Hu (2012) estimate statistical models for 44 demographic groups (based on age, gender, and educational attainment), incorporating birth cohort effects and other controls, and show that only one-fourth of the decline in the LFPR since 2008 was attributable to demographic factors. Using a multivariate Beveridge–Nelson decomposition, Van Zandweghe (2012) finds that cyclical factors accounted for 50–90% of the decline in the LFPR, depending on which measure of unemployment was used in constructing the filter. Sherk (2012) analyzes microdata from the Current Population Survey (CPS) and finds that demographic factors only accounted for one-fifth of the postrecession decline in LFPR. Finally, Hotchkiss and Rios-Avila (2013) estimate a behavioral model of labor supply using CPS microdata and conclude that the decline in LFPR since the Great Recession was more than fully explained by the deterioration in labor market conditions.[4]

We develop a simple extension of the workhorse New Keynesian model (e.g., Woodford 2003) that can account qualitatively for the stylized facts that: (i) decreases in labor force participation appear relatively modest in most postwar recessions, but (ii) protracted recessions may eventually induce large declines in participation. Our model implies that labor force participation responds inversely to the unemployment rate, but that the response is gradual due to high adjustment costs of moving between the market and "home production" sectors. Our model formulation is consistent with regression analysis we perform over the 1960–2013 period, which indicates that the LFPR responds in a very inertial manner to the unemployment rate. Given this inertial behavior, short-lived recessions have small effects on the LFPR, consistent with most postwar business cycles, while a protracted period of high unemployment can eventually cause a substantial decline in the LFPR, as occurred following the Great Recession. Importantly, to the extent that labor force participation responds very gradually to the unemployment rate, labor force participation may remain well

---

3. The unemployment gap is defined as the difference between the unemployment rate and its longer run normal rate.

4. See Hotchkiss et al. (2012), Kudlyak and Schwartzman (2012), and Daly et al. (2012) for further analysis and discussion.

below trend even as the economy begins recovering and the unemployment gap closes.

Our model also implies that the labor force participation gap enters the Phillips Curve in addition to the unemployment gap. A large negative participation gap induces labor force participants to reduce their wage demands, although our calibration implies that the participation gap has less influence than the unemployment rate quantitatively.[5] An important implication of this modified Phillips Curve is that inflation would remain below baseline following a recession even after the unemployment gap closes, at least while the participation gap remains negative.[6]

The possibility that deep recessions may generate large cyclical swings in labor force participation raises important questions for monetary policy design: in particular, what are the consequences of responding to the cyclical component of labor force participation, rather than exclusively to the unemployment rate? To address this, we use our model to analyze the implications of alternative monetary policy strategies against the backdrop of a deep recession that leaves the LFPR well below its longer run potential level. Specifically, we compare a noninertial Taylor rule, which responds to inflation and the unemployment gap to an augmented rule that also responds to the participation gap. In the simulations, the zero lower bound precludes the central bank from lowering policy rates enough to offset the aggregate demand shock for some time, producing a deep recession; once the shock dies away sufficiently, policy responds according to the Taylor rule.

A key result of our analysis is that monetary policy can induce a more rapid closure of the participation gap through allowing the unemployment rate to fall below its long-run natural rate. Quite intuitively, keeping unemployment persistently low draws cyclical nonparticipants back into labor force more quickly. Given that the cyclical nonparticipants exert some downward pressure on inflation, some undershooting of the long-run natural rate actually turns out to be consistent with keeping inflation stable in our model. However, a more aggressive strategy of employment gap targeting boosts inflation—at least to some degree—by requiring unemployment to remain lower for even longer. Thus, there is some trade-off between stabilizing inflation and broad measures of resource slack that include participation.

Policy rules that respond to broad measures of labor market slack share some characteristics of optimal *full commitment* policy strategies insofar as both imply some undershooting of the unemployment rate and overshooting of inflation as the economy recovers.[7] Even so, we stress that the undershooting of the unemployment rate in our

---

5. The implication that individuals who have left the labor force exert relatively little influence on wages is similar in spirit to that of the "insider–outsider" models of Blanchard and Summers (1986) and Blanchard and Diamond (1994). In the latter class of models, the long-term unemployed have a comparatively small influence on wage bargaining, reflecting that firms would prefer to hire workers with more recent work experience and that the long-term unemployed search less intensively.

6. Equivalently, a deep recession causes a *fall* in the short-run natural rate of unemployment, since unemployment must fall below its long-run natural rate to offset the deflationary pressure associated with the participation gap.

7. See Eggertsson and Woodford (2003) for a detailed characterization of optimal policy in the workhorse New Keynesian model under a binding zero bound constraint. The analysis of monetary

analysis occurs under a *noninertial* Taylor rule, and reflects that the participation gap remains sizably negative even after the unemployment gap (which recovers faster) has closed. By contrast, a policy of strictly stabilizing the unemployment rate generates monotonic convergence of the unemployment gap, and implies a much slower recovery in labor force participation and below-target inflation.

Overall, we view our paper as pointing out how labor market slack arising in the wake of deep recessions may not be well summarized by the unemployment rate (given the substantial lag in the response of participation). Consequentially, it may be desirable to consider adapting monetary rules developed for the Great Moderation to incorporate broader measures of slack. Of course, as we emphasize later, our model does not include a number of factors that may influence the trade-off between employment gap and inflation stabilization, and hence the results of our analysis should be interpreted very cautiously in contemplating the practical design of monetary policy.

The remainder of this paper is organized as follows. Section 1 presents our empirical analysis. Section 2 describes the specification and calibration of our model, and also characterizes the dynamics of the LFPR under a benchmark monetary policy rule. Section 3 analyzes the sources of monetary policy trade-offs, gauges the performance of simple rules and commitment-based strategies, and considers the implications of uncertainty and risks. Section 4 concludes.

1. EMPIRICAL ANALYSIS

In this section, we begin by examining the behavior of the LFPR prior to the Great Recession, and then we focus more specifically on analyzing the sources of its post-2007 decline, based on disaggregated demographic patterns and evidence from state-level panel data. We also assess the extent to which the fall in the LFPR may reflect accelerated enrollment in the Social Security Disability Insurance (SSDI) program. We then gauge the relative contribution of the participation gap and the unemployment gap to the employment gap, and finally consider how our evidence in favor of pronounced cylicality in the LFPR in the post-2007 period can be squared with the very modest degree of cyclicality in the LFPR evident through most of the postwar period.

1.1 Demographic Factors

Labor economists have long been aware of the pitfalls of using aggregate data to characterize the behavior of labor supply, not only because those characteristics can differ so markedly across demographic groups but because the magnitude of such differences can change so dramatically over time. Such considerations are clearly evident in Figure 2, which depicts the evolution of the LFPR for specific

---

strategies in our paper is positive in nature, and does not consider optimal strategies derived from maximizing a utility-based welfare function.
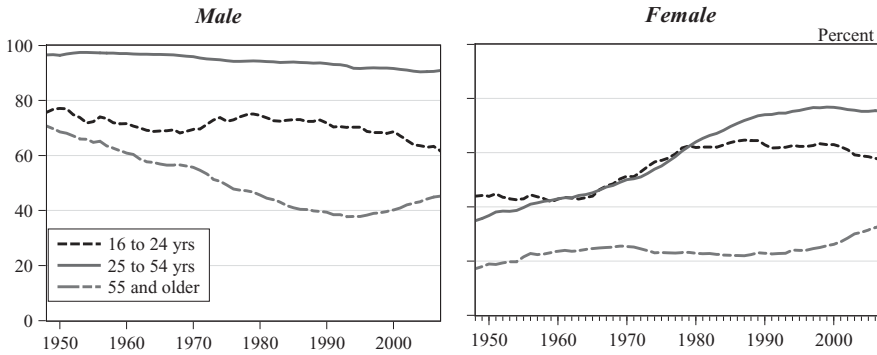
Fig. 2. Demographic Trends in Labor Force Participation.

NOTE: This figure depicts annual data regarding the labor force as a share of the civilian noninstitutionalized population (in percent) for each of the specified demographic groups over the period from 1948 to 2007.

demographic groups over the period from 1948 to 2007.[8] Of course, many volumes have been written about postwar trends in U.S. labor supply.[9] Therefore, we will simply highlight a few broad features that are salient for our empirical analysis.

Prime-age males (25–54 years) comprised about 37% of the labor force in 2007. The LFPR for prime-age males declined very gradually—about a tenth of a percentage point per year—from the late 1940s through the early 2000s. Expansions in the SSDI program account for a substantial portion of that decline, because most individuals who start receiving disability benefits never reenter the labor force; increased incarceration rates also appear to have played a significant role (see Leonard 1979, Juhn 1992, Gruber 2000, Bound and Waidmann 2002, Autor and Duggan 2003, 2006, Moffitt 2012). However, those trends appear to have subsided over the half-decade prior to the Great Recession; that is, the LFPR for prime-age males was stable at around 90.5% from 2003 to 2007.

Prime-age females (25–54 years) comprised about 31% of the labor force in 2007. The LFPR for prime-age females picked up gradually during the 1950s and 1960s, accelerated during the 1970s and 1980s, and then flattened out at a plateau of around 75%—more than twice as high as in 1948. Interestingly, microdata indicate that the wage and income elasticities of labor force participation for married females also dropped markedly over the postwar period, reaching levels that are broadly similar to those of prime-age males (see Goldin 2006, Heim 2007, and references therein.)

Youths (aged 16–24 years) comprised about 15% of the labor force in 2007. The LFPR for male youths has been on a fairly steep downward trend since the 1970s, primarily reflecting increasing rates of enrollment in postsecondary education. The

8. Figure A1 in the Appendix depicts the postwar evolution of the composition of the labor force by age and gender.

9. See Hotchkiss (2005), DiCecio et al. (2008), Autor (2010), Macunovich (2010), and Moffitt (2012) for analysis and discussion of trends in labor force participation prior to the onset of the Great Recession.

LFPR for female youths generally tracked that of prime-age females during the 1960s and 1970s but then flattened out and eventually started trending downward, moving roughly in parallel with the LFPR for male youths (see Smith 2011, and references therein).

Older adults (aged 55 years and above) comprised 17% of the labor force in 2007. The LFPR for older males declined gradually over the course of four decades to a trough of around 40% in the mid-1990s, while the LFPR for older females picked up slightly in the 1950s and then remained fairly steady at around 25% through the mid-1990s. Since then, the LFPR for older adults—both male and female—has been trending upward, apparently reflecting ongoing improvements in their overall health and ability to continue working even into the so-called "golden years."

In our view, the labor force projections published by the BLS in November 2007 serve as an invaluable resource in assessing the influence of such demographic factors on the post-2007 decline in the LFPR. In making such projections, BLS staff consider detailed demographic groups based on microdata from the CPS and other sources, including interim updates from the U.S. Census Bureau. Moreover, the timing of the November 2007 projections seems virtually perfect in terms of accomplishing our objective. At that point, most forecasters anticipated a further continuation of the Great Moderation and hence that the macroeconomy would simply move along its balanced-growth path over subsequent years. Consequently, the projected path for the labor force was closely linked to demographic factors that tend to be inertial and predictable. In retrospect, of course, the NBER dated the Great Recession as having begun just 1 month later, and hence those BLS projections effectively encompassed all of the prerecession data.

As of November 2007, the BLS projected that the aggregate LFPR would decline modestly (about 0.3 percentage point) over the half-decade from 2007 to 2012. That outlook reflected two key demographic trends, namely, the aging of the U.S. population, and the ongoing rise in the labor force participation of older adults. Indeed, in the article by Toossi (2007) in which these BLS projections were presented and discussed, the subtitle effectively captured both of those trends: *More Workers in their Golden Years*.

Regarding the first key factor, the BLS projected substantial changes in the age composition of the civilian noninstitutionalized population aged 16 years and above, as shown in Table 1. In particular, by 2012 the shares corresponding to youths and prime-age adults were expected to shrink by about 1 and 2 percentage points, respectively, while the share of older adults were projected to rise accordingly. Thus, if the LFPR for each age category had been expected to remain constant at 2007 levels, those aging patterns would have implied a downward shift of about a percentage point in the aggregate LFPR.

Regarding the second key factor, the BLS projected that the participation rates of older adults would continue rising notably over coming years, consistent with the trends that had prevailed since the mid-1990s (as noted in the discussion of Figure 2). Specifically, the LFPR for older adults (aged 55 and above) was projected to rise 2 percentage points by 2012. As for other major demographic groups,

TABLE 1

DEMOGRAPHIC FACTORS AND THE RECENT EVOLUTION OF THE LFPR

Demographic group	Population share			Labor force participation rate		
	2007 Actual	Change (2007–12)		2007 Actual	Change (2007–12)	
		Projection	Actual		Projection	Actual
16–24 years	16.1	−0.9	−0.2	59.4	−0.9	−4.5
25–54 years	54.2	−2.0	−3.1	83.0	0.3	−1.5
55–64 years	14.0	1.3	1.8	64.0	1.5	0.5
65 and older	15.6	1.7	1.6	15.4	2.6	3.1
Total	100	0	0	66.1	−0.3	−2.4

NOTES: The columns labeled "Projection" refer to the BLS labor force projections published in November 2007.

the BLS projected that the LFPR for youth would continue trending downward, while the LFPR for prime-age females would edge upward modestly and that of prime-age males would remain essentially flat—again, consistent with our discussion of the patterns shown in Figure 2.[10] All else equal, these trends in age-specific LFPRs would have pushed up the aggregate LFPR by about 0.7 percentage points.

Of course, even demographic patterns are not perfectly predictable. Table 1 shows that the actual 5-year changes in the population shares for several broad age groups have turned out to be noticeably different from what was projected in 2007; that is, the share of prime-age adults is a full percentage point lower than expected, while the shares for youths and for 55- to 64-year-olds are correspondingly higher. Nonetheless, these revisions only have modest implications for the aggregate LFPR, because the revised population shares essentially just reinforce the influence of the two key factors noted earlier. Specifically, using the actual 2012 population shares to reweight the age-specific LFPR projections implies a decline of about 0.6 percentage points in the aggregate LFPR, which is still only a fourth as large as its actual decline of 2.5%.

By contrast, as evident from the final columns of Table 1, the forecast errors for key age-specific LFPRs are large and systematic. In particular, the decline in the LFPR for youths was much steeper than the BLS had projected in November 2007, and the LFPR for prime-age adults dropped markedly rather than edging up slightly as expected. Meanwhile, the LFPR for older adults rose roughly in line with its projected path.

1.2 State-Level Data

The preceding analysis underscores a crucial question: why did the LFPR for prime-age adults decline by nearly 2 percentage points over the period from 2008 to 2012, given that the rate for this demographic group had been essentially stable over the

10. See Figure A2 in the Appendix.

preceding half-decade? In principle, that development might reflect some exogenous change in the labor–leisure preferences of prime-age adults—effectively comprising an unanticipated downward shift in their *labor supply*. The notable alternative is that this shift in prime-age LFPR was *not* a mere coincidence but instead was caused by the Great Recession and its aftermath; that is, prime-age adults dropped out of the labor force as a consequence of a large and persistent shortfall in *labor demand*.

As with many other empirical issues, the diversity of experiences across U.S. states turns out to be highly informative for distinguishing between these two hypotheses.[11] For example, the Great Recession had modest effects on the economic activity of the rural states and fairly moderate effects on certain states such as Massachusetts and Minnesota. In contrast, many other states experienced practically catastrophic outcomes: during 2008–09, the unemployment rate for prime-age adults rose more than 6 percentage points in Arizona, California, and Florida and nearly 10 percentage points in Nevada. If the post-2007 decline in prime-age LFPR simply reflected an exogenous shift in labor supply, we would expect that decline to exhibit a roughly uniform pattern apart from essentially random cross-sectional variation. On the other hand, if the drop in prime-age LFPR is indeed linked to shortfalls in labor demand, then we would expect the variation in outcomes across states to be systematically related to the cross-sectional distribution of changes in prime-age unemployment rates.

Thus, to gauge the relative important of cyclical versus structural factors, we estimate the following linear regression using ordinary least squares:

$$\Delta LFPR_i = \alpha + \beta \Delta UNEMP_i + \varepsilon_i, \tag{1}$$

where $\Delta LFPR_i$ denotes the change in the LFPR for prime-age adults in state i over the period 2007–12, and $\Delta UNEMP_i$ denotes the change in the unemployment rate of prime-age adults in that state over the period from 2007 to 2010; each series is constructed using annual average data and is measured in percentage points. In this formulation, the slope coefficient β captures the extent to which the state-level variations in prime-age LFPR tend to be associated with changes in prime-age unemployment, while the intercept α captures the extent to which prime-age LFPR exhibited a general decline across states that was unrelated to the evolution of prime-age unemployment. The results of this regression are shown in the first column of Table 2, with heteroskedasticity-consistent standard errors shown in parentheses below each coefficient estimate.

These regression results provide stark evidence that cyclical factors have been crucial in explaining the recent decline in prime-age LFPR. The coefficient on the lagged change in prime-age unemployment is highly significant (*t*-statistic of -3.9); that is, the state-level data exhibit a strong negative correlation between changes

---

11. Clark and Summers (1982) explore the relationship between the levels of LFPR and unemployment using cross-sectional data from the 1970 U.S. Census. Aaronson et al. (2006) analyze a panel of annual data for U.S. states over the period 1990–2005 and found a statistically significant relationship between the LFPR and the contemporaneous unemployment rate, using lagged unemployment rates as instruments.

TABLE 2

REGRESSION ANALYSIS OF STATE-LEVEL DATA FOR PRIME-AGE ADULTS

Explanatory variable	All U.S. states (Including Washington, DC)			41 larger states (Excluding Nevada)		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	−0.40 (0.42)	−0.41 (0.42)	—	−0.18 (0.55)	−0.18 (0.56)	—
$\Delta Unemp$ (2007–10)	−0.30 (0.08)	−0.27 (0.11)	−0.39 (0.03)	−0.33 (0.10)	−0.31 (0.14)	−0.37 (0.03)
$\Delta Unemp$ (2010–12)	—	0.10 (0.24)	—	—	0.08 (0.26)	—

NOTES: In each regression, the dependent variable is the change in the LFPR for prime-age adults in a given state over the period 2007 to 2012. Each regression equation includes the change in the annual average unemployment for prime-age adults in that state over the period from 2007 to 2010. The regressions in columns 1, 2, 4, and 5 also include a regression intercept. The regressions in columns 2 and 5 also include the change in the prime-age unemployment rate from 2010 to 2012. Standard errors are adjusted for heteroskedasticity and are given in parentheses.

in LFPR and lagged changes in unemployment for prime-age adults. In contrast, the regression intercept is not statistically significant from zero (*t*-statistic of −0.97), indicating that the data provide no support whatsoever for structural interpretations of the drop in prime-age LFPR. Given the regression coefficient of about 0.3 and that the unemployment rose about 5 percentage points nationally during the Great Recession, our state-level regressions imply a subsequent drop of about 1.5 percentage points in the national average LFPR for prime-age adults. In effect, the state-level data indicate that the aggregate decline in prime-age LFPR over the period from 2007 to 2012 can be *fully* explained by the persistent shortfall in labor demand.

Of course, before reaching any definitive conclusions, it is essential to consider the extent to which these results are robust to alternative specifications of the regression equation.[12] For example, the second column of Table 2 shows that the coefficient estimates from our benchmark regression are practically invariant to the inclusion of an additional explanatory variable, namely, the change in prime-age unemployment over the period from 2010 to 2012; moreover, the coefficient on that additional variable is negligible in size and statistically insignificant. These results indicate that the decline in prime-age LFPR from 2007 to 2012 was relatively greater in those states that experienced the largest increases in prime-age unemployment during the Great Recession and its immediate aftermath but has not been significantly affected by the relative pace of recovery across states since 2010. Indeed, that finding is broadly consistent with our characterization of the prime-age LFPR as highly inertial—a characteristic that will play a key role in the formulation of the dynamic stochastic general equilibrium (DSGE) model described in Section 2.

---

12. In addition to the results reported in Table 2, we have also confirmed that the results are broadly robust to a wide range of alternative specifications and estimation methods, including alternative definitions of the explanatory variables (e.g., timespans ending in 2009 or 2011), weighting of the observations by population size (either total or prime-age), and panel estimation with time- and state-specific effects.
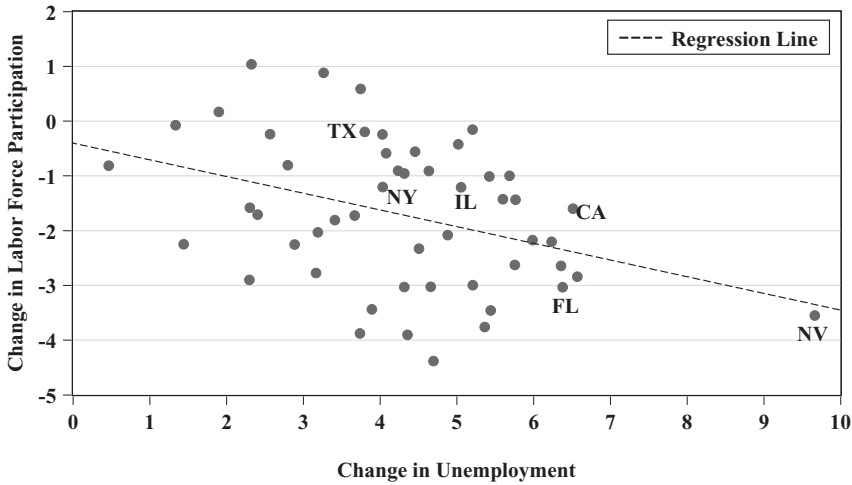
Fɪɢ. 3. State-Level Data on Unemployment and LFPR for Prime-Age Adults.

Nᴏᴛᴇs: The vertical axis refers to the change in the labor force participation rate for prime-age adults (25–54 years) between 2007 and 2012, and the horizontal axis refers to the change in the unemployment rate for that demographic group between 2007 and 2010. The dots refer to each of the U.S. states and Washington, DC. Labels annotate the observations for Nevada and the five largest states by population (California, Florida, Illinois, New York, and Texas). The dashed line depicts the regression results described in the text.

We have also confirmed that the regression results are robust to excluding potentially influential observations from the sample. In particular, as evident from Figure 3, the observation for Nevada stands out dramatically from the rest of the points in the scatterplot, and hence it is important to determine whether the coefficient estimates are affected by that particular observation. Moreover, since our benchmark regression places equal weight on all U.S. states (including Washington, DC), one might wonder whether those results are sensitive to the exclusion of states with relatively small populations. Thus, the right half of Table 2 reports the results for the sample of 41 largest states (excluding Nevada) for which the prime-age population in 2007 exceeded 500,000 people.[13] Notably, these results are virtually identical to those reported in the left half of the table.

Finally, we have recently begun to investigate these issues using a panel of state-level data on prime-age LFPRs and unemployment rates over the period from 1999 to 2012—that is, encompassing the previous business cycle as well as the more recent one. Although a full description of such analysis would go well beyond the scope of this paper, it should be noted that those results are broadly consistent with the findings reported here.

---

13. The scatterplot for that sample is shown in Figure A3 in the Appendix.

1.3 The Influence of the SSDI Program

To the extent that persistent shortfalls in labor demand have caused a large number of adults to leave the labor force, what is the likelihood that those individuals will remain permanently out of the labor force versus reengaging in the labor market once jobs become more readily available? Most notably, past empirical studies have shown that individuals who become SSDI beneficiaries almost never choose to reenter the labor force. Thus, it seems useful to consider the extent to which cyclical factors may have influenced the size of the SSDI program over the past few years.

Prime-age adults (25–54 years). As shown in the upper-left panel of Figure 4, the number of prime-age workers in the SSDI program (expressed in proportion to the civilian noninstitutionalized population of this age group) was trending upward gradually over the half-decade from 2003 to 2007 and then accelerated noticeably. By the end of 2011, the incidence of permanent disabilities was about 0.4 percentage points higher than in 2007 and about 0.2 percentage points higher than one might have predicted based on its precrisis trend. As evident from the lower-left panel, the variations across U.S. states are essentially uncorrelated with changes in state-level unemployment, suggesting that the increase in the number of prime-age SSDI beneficiaries may have primarily reflected changes in the screening process.[14]

Older adults (55 years and above). The upper-right panel of Figure 4 documents the recent acceleration in the number of workers aged 55–64 years old who have become SSDI beneficiaries. Moreover, as shown in the lower-right panel, state-level variations in unemployment exhibit a strong positive correlation with state-level changes in the number of SSDI beneficiaries in this age group, suggesting that cyclical conditions may have been quite significant in affecting workers' decisions to apply for SSDI rather than remain in the labor force. By contrast, as noted earlier, the LFPR for adults aged 65 years and older has risen notably over the past few years to a level in 2012 that was about a half percentage point *higher* than had been projected based on its precrisis trend—a development that seems consistent with widespread anecdotes about individuals who postponed their retirement in response to the drop in housing and financial wealth that was associated with the Great Recession. In effect, the average LFPR for all adults 55 years and above remains roughly in line with its precrisis trend, at least partly because the increased incidence of SSDI beneficiaries aged 55–64 has been offset by the greater incidence of working adults aged 65 and above.

Youths (16–24 years). Finally, it should be noted that there has been practically no change in the incidence of SSDI beneficiaries under age 25.[15]

14. As shown in the figure, the regression line fitted to these data is nearly horizontal, reflecting the fact that the slope coefficient is close to zero (and not statistically significant).

15. According to the SSA's latest annual report, there were 167,651 SSDI beneficiaries who were aged 16–24 years old as of December 2011, an increase of about 34,000 since December 2007—a
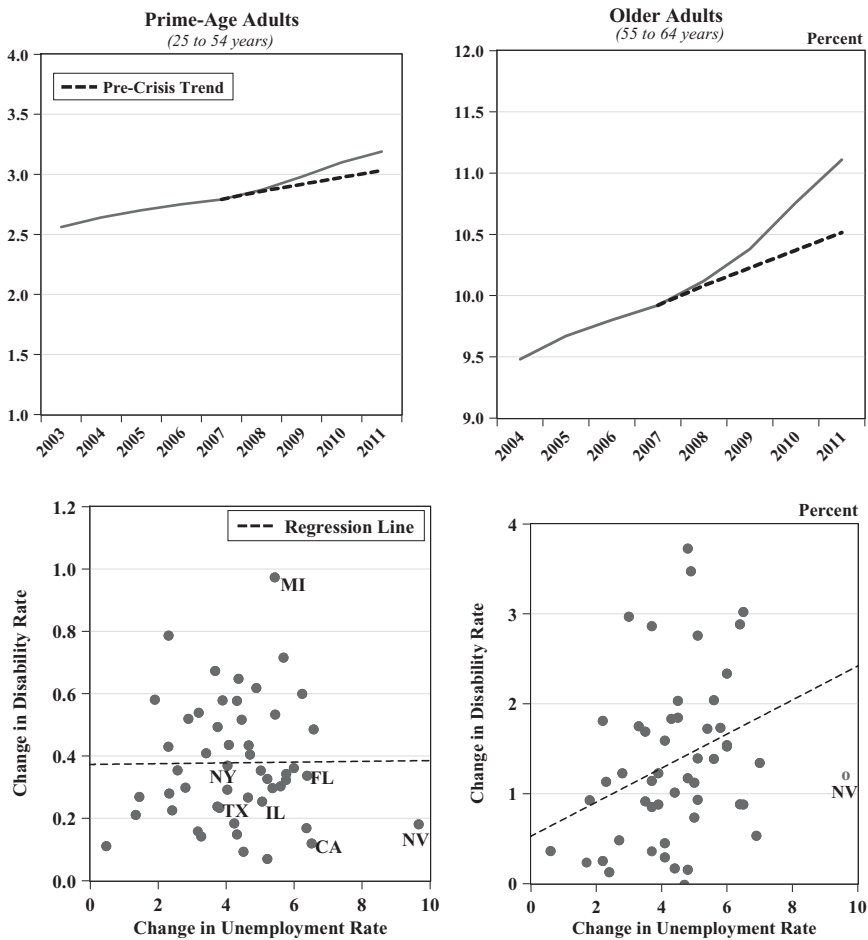
FIG. 4. The Incidence of Permanent Disabilities.

NOTES: In the upper-left panel, the solid line depicts the number of prime-age workers receiving Social Security disability benefits (i.e., beneficiaries aged 25–54 years, not including widowers or adult children) during December of each calendar year from 2003 through 2011, expressed as a percentage of the civilian noninstitutionalized population aged 25–54 years, and the dashed line denotes the projected values from fitting a linear trend over the period 2003–07. The upper-right panel provides corresponding information for older adults, that is, the number of disabled workers 55–64 years old as a share of the civilian noninstitutionalized population for that age group. In each of the lower panels, the vertical axis refers to the change in the disability rate for that demographic group between 2007 and 2011, the horizontal axis refers to the change in the unemployment rate for that group between 2007 and 2010, and the dashed line depicts the regression results as described in the text.

miniscule change relative to the total population of about 38 million in this age group. This total number of SSDI beneficiaries includes those designated as "workers" (whose eligibility is linked to their own work experience prior to becoming disabled) as well as those designated as "adult children" (who became

1.4 Gauging the Magnitude of the Employment Gap

We now proceed to consider the implications of the foregoing discussion for assessing the magnitude of the employment gap ($EGAP_t$), that is, the percentage point deviation between civilian employment (E_t) and its "natural" (or long-run sustainable) level (E_t^*) as a share of the civilian noninstitutionalized population (POP_t). In particular, the employment gap satisfies the following relationship:

$$EGAP_t = (1 - u_t^*)(LFPR_t - LFPR_t^*) - LFPR_t^*(u_t - u_t^*)$$
$$- (u - u_t^*)(LFPR_t^* - LFPR_t), \tag{2}$$

where $LFPR_t$ refers to the actual LFPR, $LFPR_t^*$ denotes its *structural* rate based solely on demographic trends, u_t refers to the unemployment rate, and u_t^* denotes the natural rate of unemployment (or NAIRU). The natural employment rate E_t^* is simply defined as $E_t^* = LFPR_t^*(1 - u_t^*)$. To facilitate the discussion, it is convenient to define the participation gap as $LFPR_t - LFPR_t^*$, while the unemployment gap is given by $(u_t - u_t^*)$. The final term in equation (2) is simply the product of these two gaps and is generally negligible compared to the preceding terms in that equation. Consequently, the employment gap can be closely approximated as the weighted sum of the participation gap and the unemployment gap:

$$EGAP_t \simeq (1 - u_t^*)(LFPR_t - LFPR_t^*) - LFPR_t^*(u_t - u_t^*). \tag{3}$$

Thus, the employment gap becomes more negative as the participation rate declines relative to its natural rate, or if the unemployment rate rises relative to the NAIRU. The two gaps are multiplied by the proportionality factors $(1 - u_t^*)$ and $LFPR_t^*$, respectively. In effect, the weight on the participation gap is slightly below unity (say, 0.95), whereas the weight on the unemployment gap is substantially smaller (say, 0.67) because the incidence of unemployment is measured as a fraction of the labor force whereas the participation rate is constructed in terms of the civilian noninstitutionalized population. To operationalize this formula, we take the trend path $LFPR_t^*$ as projected by the BLS in November 2007, and we use the CBO's 10-year-ahead unemployment rate forecast as the value of u_t^*.

As shown in the left panel of Figure 5, the employment gap widened markedly during the Great Recession to a trough of about -4% by mid-2009, corresponding to a shortfall of nearly 10 million jobs. At that point, the unemployment gap accounted for the bulk of the employment gap, while labor force participation remained roughly in line with its precrisis trend. This panel also underscores the excruciatingly sluggish nature of the recovery. The employment gap has only narrowed modestly over the past few years, because the steady widening of the participation gap has been roughly

permanently disabled prior to age 22 and whose eligibility was based on the work experience of the person's parents).
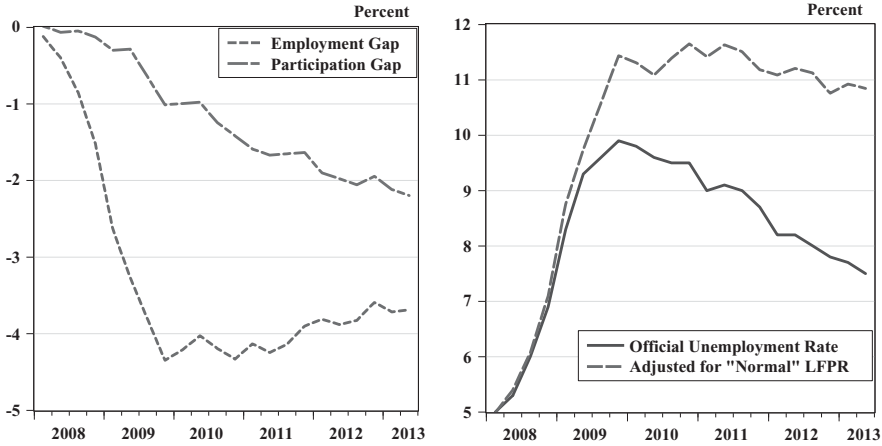
FIG. 5. The Magnitude and Composition of the Employment Gap.

NOTES: In the left panel, the employment gap refers to the percent deviation between U.S. civilian employment and its natural level (solid line), and the participation gap refers to the deviation of the LFPR from the CBO's assessment of its potential path (dashed line), where the latter is scaled by the proportionality factor $(1−u^*)$. In the right panel, the "official" unemployment rate (solid line) refers to the series U3 as published by the BLS, and the unemployment rate under "normal" LFPR (dashed line) refers to the sum of the unemployment rate and the labor force gap.

comparable in magnitude to the decline in the unemployment gap.[16] Indeed, it is evident that as of 2013:Q3, the participation gap exceeded 2 percentage points, and accounted for more than half of the employment gap. The gradual widening of the participation gap is consistent with evidence from state-level data of a protracted response of the LFPR to employment conditions.[17]

An alternative way of representing the employment gap is to consider what level of unemployment would prevail if those individuals who dropped out of the labor force due to cyclical factors had instead continued searching actively for work. In particular, we can construct an "adjusted" measure of the unemployment rate U_t^a that would prevail if the LFPR followed its "normal" trajectory based solely on structural factors:

$$U_t^a = U_t + (1 - u_t^*)(LFPR_t - LFPR_t^*)/LFPR_t^*, \tag{4}$$

where U_t denotes the conventional measure of unemployment (which the BLS refers to as U3). By comparing equations (3) and (4), it is evident that the

16. The BLS classifies individuals as "marginally attached" to the labor force if they have searched for a job within the past year (but not within the past month) and indicate that they would like a job and are available to work at the present time. As shown in Figure A4 in the Appendix, such individuals only comprise a modest fraction of the overall participation gap.

17. While the empirical analysis in this paper focuses on the United States, Howard, Martin, and Wilson (2011) find that recessions in industrial economies over the 1960-2010 period were usually accompanied by a gradual but eventually sizable drop in the LFPR (based on an event study analysis of 46 recession episodes). The LFPR typically remained well below its precrisis trend even 5 years after the onset of recession.

adjusted unemployment rate conveys the same information as the employment gap, but the former is gauged in terms of the structural path of the labor force whereas the latter is gauged in terms of the civilian noninstitutionalized population.

As shown in the right panel of Figure 5, the adjusted measure of unemployment climbed to around 11.5% by the end of 2009—about 1.5 percentage points higher than the conventional measure of unemployment. The adjusted unemployment rate remained close to its postrecession peak through mid-2011 and subsequently declined only modestly. In effect, the conventional measure of unemployment appears to have significantly overstated the actual extent of labor market recovery in the wake of the Great Recession.

1.5 Evidence from Past Business Cycles

An important question is whether the large cyclical decline in the LFPR that occurred in the wake of the Great Recession represents a distinct break from the pattern evident in previous postwar business cycles. To help address it, we use CBO estimates of the "potential" level of the labor force E_t^* to derive a labor force gap measure $LFGAP_t$ defined as the percent difference between the labor force and its potential level; see CBO (2013). Assuming that business-cycle-related fluctuations in the working age population are very small, the participation gap is essentially proportional to the labor force gap up to a scale factor equal to the steady-state LFPR, that is, $LFPR_t - LFPR_t^* = LFPR(LFGAP_t)$. Thus, given that the CBO series extends back to 1960, we can assess cyclicality in the participation gap over a long postwar sample.

The CBO potential labor force measure is appealing insofar as it is based on statistical analysis of highly disaggregated data. However, it is important to point out a key difference between our preferred measure of the potential LFPR $LFPR_t^*$ that was used in Figure 5 and the CBO measure. In particular, the CBO measure of potential is an *ex post* measure that is affected by any business-cycle developments that are expected to have enduring effects on the LFPR. Thus, adverse developments associated with the Great Recession that induced people to enroll in the SSDI program, or that persistently weakened labor force attachment, would reduce the CBO measure of the potential LFPR. By contrast, our preferred measure of the $LFPR_t^*$ is based on an *ex ante* forecast, and hence abstracts from any such hysteresis effects on potential; this reflects that we would like to capture such hysteresis effects in our measure of the participation gap. As a result, the CBO-based participation gap following the Great Recession is somewhat smaller than our preferred measure (as we will next show). Even so, this "downward bias" in the CBO measure of the participation gap (in absolute terms) is likely to be relatively small prior to 2007 given that unemployment spells were typically much shorter in duration.

Figure 6 depicts the evolution of the CBO's labor force gap $LFGAP_t$ and the unemployment gap over the past few decades.[18] Although the labor force gap shows

---

18. Given that the CBO data on the potential labor force are published at an annual frequency, we use Eviews' quadratic interpolation method in constructing our quarterly time series.
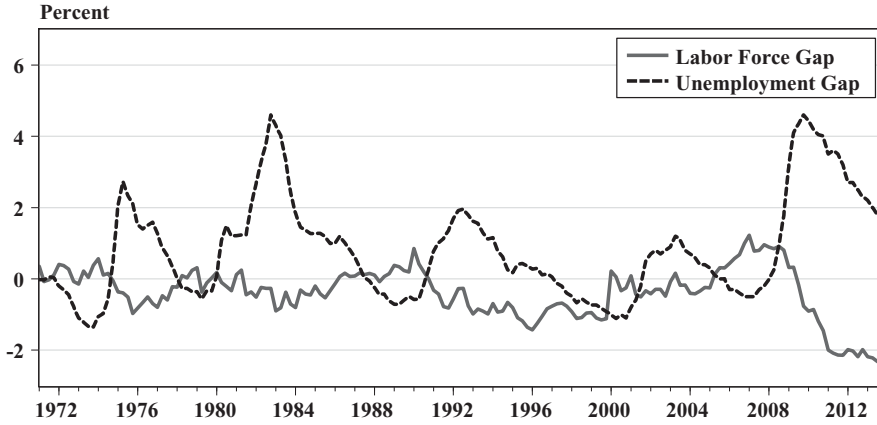
Fɪɢ. 6. The Labor Force Gap and the Unemployment Gap.

NOTES: The labor force gap is the percent deviation between the actual labor force and the CBO's assessment of its potential level. The unemployment gap is the deviation between the actual unemployment rate and the CBO's assessment of the longer run NAIRU.

clear business-cycle variation, the magnitude of the fluctuations was quite modest prior to 2007, especially compared to the large swings in unemployment that occurred during the mid-1970s and the early 1980s. In effect, the pre-2007 data seem reasonably consistent with the conventional wisdom that the LFPR could be characterized as practically acyclical. By contrast, in the wake of the Great Recession, the labor force gap has widened to an unprecedented degree. The 2% labor force gap implies an increase in the LFPR gap of about 1.5 percentage points, only modestly smaller than the 2 percentage point gap shown in Figure 5.

To gauge the degree of cyclicality more precisely, we regress the labor force gap on its own lagged value and on the lagged value of the unemployment gap:

$$LFGAP_t = \alpha LFGAP_{t-1} + \beta UGAP_{t-1} + \varepsilon_t \tag{5}$$

using quarterly data over the period 1960:Q1 through 2013:Q3.[19]

As shown in Table 3, the coefficient estimate on the contemporaneous unemployment gap is negative and highly statistically significant in the full sample, and also in subsamples determined by a 1990:Q1 breakpoint. Because the coefficient is small in magnitude, relatively short-lived increases in the unemployment gap have small effects on the LFPR. This helps account for the practically acyclical behavior of the LFPR prior to 2007, given that previous postwar recessions were fairly short-lived; indeed, even sharp recessions in the mid-1970s and early 1980s did not have much impact on the LFPR because the subsequent recoveries were V-shaped. Nonetheless,

19. We have confirmed that additional lags are not statistically significant for the entire sample or for the two subsamples reported in Table 1.

TABLE 3

Gauging the Cyclicality and Persistence of the Labor Force Gap

Sample	Persistence (α)	Cyclicality (β)
Full (1960:Q1 to 2013:Q3)	0.881 (0.027)	−0.068 (0.015)
Pre-1990 (1960:Q1 to 1989:Q4)	0.785 (0.043)	−0.066 (0.019)
Post-1990 (1990:Q1 to 2013:Q3)	0.917 (0.037)	−0.075 (0.021)

Notes: For each sample, this table reports the coefficient estimates obtained from regressing the labor force gap on its own lagged value (α) and on the lagged value of the unemployment gap (β). Standard errors are adjusted for heteroskedasticity and are given in parentheses.

because the dynamics of the LFPR are highly inertial—with a persistence coefficient exceeding 0.9 in the post-1990 subsample—our estimates imply a much larger cyclical response of the LFPR following a deep and protracted period of high unemployment, as occurred following the Great Recession. For example, while a 2 percentage point rise in the unemployment gap lasting only a year would cause the LFPR to fall only about 0.3%, a rise in the unemployment gap calibrated to match that which occurred in the Great Recession would induce a peak decline in the participation rate of 1.7 percentage points—just a bit less than the 2 percentage point drop shown in Figure 5.

Overall, these results suggest that behavior of the LFPR post-2007 seems broadly in line with the sensitivity of the LFPR to the unemployment gap over a much longer stretch of the postwar period. Although it is plausible that the responsiveness of the LFPR to unemployment may have increased slightly in recent years, the main reason that the LFPR has declined so markedly since 2007 is the deep recession and unusually sluggish nature of the recovery in employment. Of course, while the foregoing analysis is suggestive, much more work is needed to better gauge how the LFPR is likely to respond to labor market conditions—including consideration of more disaggregated measures of labor force participation, and alternative explanatory variables such as short- versus long-term unemployment.

2. MODEL SPECIFICATION

In this section we describe a New Keynesian model with endogenous labor force participation. Our formulation of the participation decision draws on the literature that incorporated home production into macromodels (e.g., Benhabib, Rogerson, and Wright 1991, and Hercowitz 1991). In the spirit of the home production literature, labor supply decisions are made by a representative household that chooses to allocate labor between the market and home (or nonmarket) sectors based on the relative return to working in either sector. However, while the home production literature focused on how relative productivity differentials between market and nonmarket activities

affect labor flows, our New Keynesian framework is suited to analyzing how weak demand conditions in the market sector—associated with high unemployment and a low return to working—affect labor force participation.

A second key feature of our model is the inclusion of adjustment costs of moving labor between the market and home production sectors (with the home production sector corresponding to individuals who have left the labor force). Adjustment costs help account for why relatively transient spells of unemployment—the experience of most postwar recessions—appear to cause only small changes in participation, while deep recessions may eventually cause large changes in the LFPR. Because the participation rate is a slowly evolving state variable, participation may remain low or even fall as the economy recovers and unemployment drops; as we highlight in the next section, this divergence between the participation rate and unemployment rate may have crucial implications for monetary policy.

Notwithstanding a number of technical assumptions described later, our framework attempts to capture in a simple way the traditional Keynesian notion that weak aggregate demand leads to unemployment, and that labor force participation is likely to respond sluggishly. On the unemployment dimension, our model in effect reinterprets the "hours gap" of the workhorse New Keynesian model along the extensive margin of unemployment.[20] An alternative approach would be to model unemployment and labor force participation in a full-scale search model; while this seems desirable in general, recent research by Veracierto (2008) and Shimer (2008) suggests that a fairly rich and complicated model structure may be required to account for empirically plausible patterns of comovement between output, unemployment, and labor force participation.[21]

2.1 Households

Our model assumes that a representative household allocates "families" between the market and home production sectors. Each family is itself composed of many members (i.e., a continuum). Assuming a continuum of such families on the unit interval, the representative household chooses a fraction L_{Ft} to work in the market sector, and $1 - L_{Ft}$ to work in the home production sector; thus, L_{Ft} is the LFPR.

This decentralization involving families with a continuum of individual members is convenient for introducing unemployment into the model. Similar to Gali (2011), we assume that individual family members can be regarded as ordered sequentially based on their disutility of working during the period, with individual h experiencing a

---

20. The workhorse New Keynesian model in fact is a special case of our model in which the LFPR is fixed.

21. For example, both Veracierto (2008) and Shimer (2008) find that the unemployment rate varies *procyclically* if prices and wages are fully flexible. Shimer finds that incorporating sticky (real) wages helps account for the correct pattern of procyclical participation and a countercyclical unemployment rate, but only provided that the disutility of unemployment is roughly commensurate with that of working (if the disutility of unemployment is considerably smaller, then unemployment is procyclical even with the wage rigidities).

disutility of work of $\chi_0 h^\chi$ if hired to work in the market sector during the period (with all workers working the same fixed number of hours), and zero disutility otherwise. Importantly, the "type" of each family member is revealed each period only after the labor force participation decision has been made, so that individual family members cannot be allocated to the market and nonmarket sector based on their individual (dis)taste for market work: the representative household moves the entire family to one sector or the other. In order to satisfy the economy-wide employment demand for $N_t = L_{Ft} H_{mt}$ workers, each of the L_{Ft} families in the market sector hires those members with the lowest disutility of work during the period. The total disutility of work to a family in the market sector is given by $V(H_{mt}) = \chi_0 \int_0^{H_{mt}} h^\chi dh = \frac{\chi_0}{1+\chi} H_{mt}^{1+\chi}$, where H_{mt} is the number of family members employed, or equivalently, H_{mt} is the employment rate of the labor force (reflecting that all families behave identically), and $1 - H_{mt}$ the unemployment rate.

The representative household, which maximizes the utility of all "families" in the economy, has a utility functional of the form:

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \epsilon_{dt+j} \left\{ \begin{array}{c} \frac{1}{1-\sigma_C} \left(C_{t+j} - \psi_c C_{t+j-1}^a - \varepsilon_c \varepsilon_{ct} \right)^{1-\sigma_C} - L_{Ft+j} \frac{\chi_0 H_{mt+j}^{1+\chi}}{1+\chi} + \frac{\gamma_G}{1-\sigma_G} L_{Ht+j}^{1-\sigma_G} \\ -0.5\phi_G \left(\frac{L_{Ft+j}}{L_{Ft+j-1}} - 1 \right)^2 L_{Ft+j} \end{array} \right\}, \quad (6)$$

where the discount factor β satisfies $0 < \beta < 1$, and ϵ_{dt} is a shock to the discount factor. All consumption, including both market consumption C_t and home production L_{Ht}, is pooled across families so that all families enjoy equal consumption (as do all individuals within a family). The period subutility function over consumption allows for the possibility of external habit persistence in consumption, with C_{t-1}^a denoting lagged aggregate consumption, and also allows for a shock to the marginal utility of consumption ε_{ct}. The period disutility of work reflects that each family allocated to market production sector experiences a disutility of $V(H_{mt}) = \frac{\chi_0}{1+\chi} H_{mt}^{1+\chi}$, so that the cumulative disutility of the L_{Ft} working families is $L_{Ft} \frac{\chi_0 H_{mt+j}^{1+\chi}}{1+\chi}$. Home production yields a period utility benefit $M(L_{Ht}) = \frac{\gamma_G}{1-\sigma_G} L_{Ht}^{1-\sigma_G}$ that rises in the number of families allocated to that sector (i.e., $L_{Ht} = 1 - L_{Ft}$) minus adjustment costs incurred from shifting resources across sectors (the last term). The adjustment costs are assumed to be quadratic, and may be either internal, or external if the lagged labor force is taken as given by the representative household.[22]

The representative household's budget constraint in period t states that its expenditure on goods and net purchases of (zero-coupon) government bonds must equal its disposable income:

$$P_t C_t + B_{G,t} - (1 + i_{t-1}) B_{G,t-1} = W_t N_t + \Gamma_t. \quad (7)$$

---

22. In our framework, worker productivity does not depend on the time spent in or outside of the labor force. Given that decisions to participate in the workforce involve weighing the effects of skill deterioration on future earnings, an interesting extension would be to allow for the effects of on-the-job training on wages (e.g., Chang, Gomez, and Schorfheide 2002).

Here $B_{G,t}$ are purchases of bonds that promise a nominal return of $(1 + i_t)$ in the following period, $W_t N_t$ is (nominal) wage income, and Γ_t is income from profits. Total employment of the representative household N_t is the product of employment rate of each household working in the market sector H_{mt} and the LFPR L_{Ft}, that is, $N_t = H_{mt} L_{Ft}$.

The first-order condition for the employment rate H_{mt} is simply given by:

$$\frac{dV(H_{mt})}{dH_{mt}} = \chi_0 H_{mt}^{\chi} = \frac{W_t}{P_t} \lambda_t. \tag{8}$$

For each working family, equation (8) defines the threshold disutility of work such that all family members with a disutility below this level choose to work, while all other members remain unemployed. This condition is exactly equivalent to that in the standard New Keynesian model, which equates the marginal disutility of working to the real wage; the only difference is that here the marginal disutility is interpreted along the extensive margin of employment, rather than the intensive margin of hours worked. Thus, a shortfall in aggregate demand is interpreted as reducing the employment rate relative to its flexible price level, rather than reducing hours worked relative to its level under flexible prices.

Abstracting from adjustment costs, the first-order condition for labor force participation L_{Ft} is given by:

$$\left\{ \lambda_t \frac{W_t}{P_t} H_{mt} - V(H_{mt}) \right\} = \frac{dM(L_{ht})}{dL_{ht}} = (1 - L_{ft})^{-\sigma_G}. \tag{9}$$

Equation (9) implies that the representative household chooses to allocate families to market production up to the point at which marginal return to market work equals the marginal cost in terms of foregone household production $\frac{dM(L_{ht})}{dL_{ht}}$. The marginal return to allocating another family to market work—in brackets on the left-hand side—equals the family's total wage income (expressed in utils) minus the total disutility $V(H_{mt})$ that the family would experience from working in the market sector. Noting that $V(H_{mt}) = \frac{dV(H_{mt})}{dH_{mt}} \frac{H_{mt}}{1+\chi}$ (reflecting the isoelastic specification of $V(H_{mt})$) and using equation (8)), $V(H_{mt})$ may be expressed as $V(H_{mt}) = \lambda_t \frac{W_t}{P_t} \frac{H_{mt}}{1+\chi}$. Substituting into equation (9) and using (8) yields:

$$\left\{ \frac{\chi}{1+\chi} \lambda_t \frac{W_t}{P_t} H_{mt} \right\} = \frac{\chi}{1+\chi} \chi_0 H_{mt}^{1+\chi} = (1 - L_{ft})^{-\sigma_G}. \tag{10}$$

Thus, absent adjustment costs, labor force participation L_{ft} varies directly with the employment rate H_{mt}. Quite intuitively, factors that increase the return to market work—and hence boost the employment rate—also increase labor force participation, so that households adjust on both margins. As we discuss later, adjustment costs slow the response of labor force participation to employment changes.

Finally, the optimal bond holding choice of the representative household implies the condition:

$$\lambda_t = E_t \left(\beta \frac{\epsilon_{dt+1}}{\epsilon_{dt}} \right) \lambda_{t+1} \frac{1 + i_t}{1 + \pi_{t+1}}, \tag{11}$$

where the marginal utility of consumption is given by $\lambda_t = (C_t - \psi_c C_{t-1}^a)^{-\sigma_c}$, and π_t is the inflation rate ($\frac{P_t}{P_{t-1}} - 1$).

2.2 Firms

On the production side, we assume a familiar setting with a continuum of monopolistically competitive firms to rationalize Calvo-style price stickiness. Each firm has a production function that depends on capital $K_t(f)$ and labor $N_t(f)$ of the form:

$$Y_t(f) = K_t(f)^\alpha N_t(f)^{1-a}. \tag{12}$$

While the aggregate capital stock is fixed, capital may be freely allocated across the f firms, implying that real marginal cost $MC_t(f)/P_t$ is identical across firms and equal to

$$\frac{MC_t}{P_t} = \frac{W_t/P_t}{MPL_t} = \frac{W_t/P_t}{(1-\alpha)K^\alpha N_t^{-\alpha}}. \tag{13}$$

Each monopolistically competitive firm faces a downward-sloping demand curve of the form $Y_t(f) = [\frac{P_t(f)}{P_t}]^{\frac{-(1+\theta_p)}{\theta_p}} Y_t$, where θ_p determines the elasticity of substitution between the differentiated goods, and equals the net markup. Given Calvo-style pricing frictions, firm f that is allowed to reoptimize its price ($P_t^*(f)$) solves the usual problem:

$$\max_{P_t^*(f)} E_t \sum_{j=0}^{\infty} \xi_p^j \psi_{t,t+j} [(1+\pi)^j P_t^*(f) - MC_{t+j}] Y_{t+j}(f),$$

where $\psi_{t,t+j}$ is the stochastic discount factor (the conditional value of future profits in utility units, that is, $\beta^j E_t \frac{\lambda_{t+j} P_{t+j}}{\lambda_t P_t}$). To allow for the possibility of structural persistence in inflation, we implement in practice a well-known variant in which a fraction ν of those firms that do not receive a signal to reoptimize mechanically adjust their price in line with past inflation.

The aggregate resource constraints for the economy imply that output equals consumption, that is, $Y_t = C_t$, and that the supply of capital and labor used by the monopolistically competitive firms sum to the relevant aggregates, that is, $\int K(f)df = K$ and $\int N_t(f)df = N_t$.

2.3 Log-Linearized Model

Given that prices are sticky, output and employment in our model are demand determined. Aggregate demand in the log-linearized version of our model can be expressed in terms of the familiar "New Keynesian" IS curve:[23]

$$n_t - n_t^* = (n_{t+1|t} - n_{t+1|t}^*) - \frac{1-\alpha}{\hat{\sigma}_c}(i_t - \pi_{t+1|t} - r_t^*). \tag{14}$$

In the foregoing equation, n_t is the employment rate and n_t^* is the "natural" or "potential" employment rate that would prevail under fully flexible prices, with each variable expressed as log percent deviation from its steady-state value.[24] We use the notation $n_{t+j|t}$ to denote the conditional expectation of a variable n at period $t+j$ based on information available at t, that is, $n_{t+j|t} = \mathrm{E}_t n_{t+j}$.

Equation (14) indicates that the the employment gap $n_t - n_t^*$ depends inversely on the deviation of the real interest rate $(i_t - \pi_{t+1|t})$ from its natural rate r_t^*, as well as on the expected employment gap in the following period. Given that output is simply proportional to employment (i.e., $y_t = (1-\alpha)n_t$), this equation may be expressed alternatively in terms of the output gap $y_t - y_t^*$. The interest sensitivity of the employment gap in equation (14) depends both on the labor share $1-\alpha$ and on the household's intertemporal elasticity of substitution in consumption $\frac{1}{\hat{\sigma}_c}$; the composite parameter $\hat{\sigma}_c$ is defined as $\hat{\sigma}_c = \frac{\sigma_c}{1-\varepsilon_c}$, and is essentially equal to the curvature parameter σ_c in the household utility function given that the mean level of the consumption taste shock ε_c is assumed to be small.

Our model implies that a first-order log approximation to the employment gap can be expressed as:

$$n_t - n_t^* = (l_{ft} - l_{ft}^*) - \frac{1}{1-u}(u_t - u_t^*). \tag{15}$$

Thus, the employment gap may turn negative due to either a negative labor force participation gap $(l_{ft} - l_{ft}^*)$, or because unemployment rises above its natural rate (i.e., a positive unemployment rate gap $u_t - u_t^*$). To a first-order approximation, equation (15) is identical to equation (3) in Section 1.4 except that the employment gap here is expressed as a (log) percent deviation of the employment rate from its natural rate, rather than as a percentage point deviation, and similarly for the participation rate. While the conventional measure of the unemployment gap is simply $u_t - u_t^*$, it is convenient to refer to the scaled version $\frac{1}{1-u}(u_t - u_t^*)$ as the unemployment gap, reflecting that the scaling factor translates changes in the participation gap and unemployment gap into commensurate units.

---

23. For expositional simplicity, we abstract from habit persistence in presenting these log-linearized equations, but we do incorporate habit persistence in the model simulations reported later.

24. More generally, the variables in the log-linearized equations are measured as percent or percentage point deviations from their steady-state level. The superscript "*" denotes the level of a variable that would prevail under completely flexible prices, which we refer to as the "natural" or potential level (e.g., u_t^* is the natural rate of unemployment).

The price-setting equation specifies the inflation rate π_t to depend both on expected inflation and marginal cost mc_t:

$$\pi_t - \nu\pi_{t-1} = \beta(\pi_{t+1|t} - \nu\pi_t) + \kappa_p mc_t. \tag{16}$$

Reflecting the Calvo–Yun contract structure, the composite parameter $\kappa_p = \frac{(1-\xi_p)(1-\beta\xi_p)}{\xi_p}$, and hence depends inversely on the mean contract duration ($\frac{1}{1-\xi_p}$). The quasi-difference specification for inflation reflects that we allow for some degree of dynamic indexation in price contracts.

A novel feature of the aggregate supply block is that marginal cost mc_t depends both on the unemployment rate gap and the labor force participation gap:

$$mc_t = \frac{w_t}{p_t} - MPL_t = \psi_u \frac{1}{1-u}(u_t^* - u_t) + \psi_l(l_{ft} - l_{ft}^*), \tag{17}$$

Equations (16) and (17) together imply that inflation falls relative to target if there is slack in *either* labor market resource gap, that is, if either $u_t > u_t^*$ or $l_{ft} < l_{ft}^*$. The implication that a negative cyclical participation gap would exert at least *some* downward pressures on wages and marginal cost seems reasonable, especially to the extent that it reflected a falloff in the participation rate of younger and geographically mobile members of the population.

The Phillips Curve implied by our model (i.e., equations (16) and (17)) clearly resembles that associated with the "insider–outsider" models of the labor market that were pioneered by Blanchard and Summers (1986) and Blanchard and Diamond (1994). An important feature of this class of models is that the longer term unemployed affect wages by less than the newly unemployed, and empirical applications attempting to differentiate between these groups in estimating price- or wage-Phillips Curves typically found support for this prediction (e.g., see Blanchard 1991). Drawing on this literature, it seems plausible that cyclical nonparticipants in our model would have relatively little influence on wages and marginal cost relative to labor force participants—a feature that we impose in our calibration (i.e., ψ_u is much larger than ψ_l).

Although our model literally imposes that the coefficient ψ_u on the unemployment gap exceed the coefficient on the participation gap ψ_l, the model can be calibrated flexibly so that the gap between coefficients approaches zero as the elasticity of labor supply becomes arbitrarily high (noting $\psi_u = \chi + \psi_l$, where $\psi_l = \alpha + (1 - \alpha)\hat{\sigma}_c$, and recalling that χ is the inverse Frisch elasticity of labor supply). This flexibility clearly seems desirable given substantial uncertainty about the relative importance of each gap in affecting marginal costs. The reason why the unemployment gap coefficient is *generally* higher in our model reflects that the participation gap only affects marginal cost through a wealth effect on wages and due to diminishing returns to production, whereas the unemployment gap—in addition to working through these channels—also affects wages through its effect on the disutility of work. For example, in the special case of a linear production function ($\alpha = 0$) and no taste shocks, the

log-linearized version of equation (8) is given by:

$$mc_t = -\frac{\chi}{1-u}(u_t - u_t^*) + \hat{\sigma}_c(c_t - c_t^*)$$

$$= -\left(\frac{\chi + \hat{\sigma}_c}{1-u}\right)(u_t - u_t^*) + \hat{\sigma}_c(l_{ft} - l_{ft}^*). \tag{18}$$

Thus, the participation gap enters only through the wealth effect term $\hat{\sigma}_c(c_t - c_t^*)$, which captures how (say) lower participation reduces consumption and hence shifts out the wage schedule.

The log-linearized equation describing how the labor force participation gap responds dynamically to the unemployment rate gap is given by:

$$l_{ft} - l_{ft}^* = \mu_1(l_{ft-1} - l_{ft-1}^*) - \phi_m(u_t - u_t^*), \tag{19}$$

under the assumption that adjustment costs are external to the representative household. The lag coefficient μ_1 captures the magnitude of adjustment costs, while the composite parameter $\phi_m = \frac{(1-\mu_1)(1+\chi)}{\sigma_G(\frac{l_f}{1-l_f})(1-u)}$ determines the contemporaneous response of participation. Quite intuitively, shocks that raise the return to working in the market sector—and thus lower the unemployment rate—also encourage a higher participation rate, and conversely for a contractionary shock. The response coefficient ϕ_m depends on the relative degree of curvature of the disutility of work schedule to that of the home production schedule: for example, as the home production function becomes more concave (higher σ_G), the responsiveness of participation to a change in the return to market work is diminished. It is worth noting that under internal adjustment costs, labor force participation would also respond to expectations about future unemployment (since a lead of participation would enter into equation (19)); while this seems an interesting extension for future work, we focus on equation (19) given that it lines up well with the empirical analysis in Section 1.5 and has the appeal of simplicity.

In our baseline model, we assume that the policy rate i_t follows a simple Taylor rule subject to the zero lower bound:

$$i_t = \max(-i, \gamma_\pi \pi_t + \gamma_n(n_t - n_t^*) - \gamma_u(u_t - u_t^*) + \epsilon_{it}). \tag{20}$$

In addition to inflation, the policy rate is assumed to react to the unemployment rate gap $u_t - u_t^*$, to the employment gap $n_t - n_t^*$, or possibly both. The policy rule is assumed to be noninertial, though it is worth noting that a policy rule reacting to the employment gap may be regarded as responding indirectly to labor force participation, which is a state variable.

Potential employment n_t^*, potential labor force participation l_{ft}^*, and the potential unemployment rate u_t^*—which are the values these variables would assume if prices

were always flexible—are determined by:

$$-\psi_u \frac{1}{1-u}u_t^* + \psi_l l_{ft}^* = \hat{\sigma}_c \varepsilon_c \varepsilon_{ct}, \tag{21}$$

$$l_{ft}^* = \mu_1 l_{ft-1}^* - \phi_m u_t^*, \tag{22}$$

$$n_t^* = l_{ft}^* - \frac{1}{1-u}u_t^*. \tag{23}$$

These equations are the analogues of equations (17), (19), and (15), respectively. The first equation (21) indicates that a contractionary taste shock ε_{ct}, which reduces the marginal utility of consumption, would raise the potential unemployment rate if the participation rate were held unchanged (reflecting an inward shift in the labor supply curve). But given that it is efficient to adjust to shocks on the participation margin as well, equation (22) indicates that the negative taste shock also eventually reduces potential labor force participation.

The natural (potential) real interest rate may be expressed as:

$$r_t^* = \hat{\sigma}_c[(1-\alpha)(n_{t+1|t}^* - n_t^*) + \varepsilon_c(\varepsilon_{ct} - \varepsilon_{ct+1|t})] + \varepsilon_d(\varepsilon_{dt} - \varepsilon_{dt+1|t}). \tag{24}$$

We assume that the consumption taste shock ε_{ct} and discount factor shock ε_{dt} follow AR(1) processes with persistence parameters of ρ_c and ρ_d, respectively. Accordingly, a negative consumption taste shock depresses the natural real interest rate. Although discount factor shocks have no effect on potential employment n_t^* or its components, a temporary discount factor shock that increases the preference for future relative to current consumption (implying $\varepsilon_{dt} < \varepsilon_{dt+1|t}$) also reduces the natural real rate.

2.4 Calibration

To run our benchmark simulations, we calibrate the model as follows. We set the discount factor $\beta = 0.995$, and the steady-state net inflation rate $\pi = 0.005$; this implies a steady-state interest rate of $i = 0.01$ (i.e., 4% at an annualized rate). We set the parameter $\hat{\sigma}_c$ determining the interest elasticity of employment demand equal to unity (with the scale parameter on the consumption taste shock $\varepsilon_c = 0.01$), the capital share parameter $\alpha = 0.3$, and the Frisch elasticity of labor supply $\frac{1}{\chi} = 0.25$. The habit persistence in consumption parameter ψ_c is set to 0.8. These parameter values imply that $\psi_l = 1$ and $\psi_u = 4$ in the marginal cost expression (17), so that a 1 percentage point rise in the participation gap boosts marginal cost by one-fifth as much as a 1 percentage point fall in the unemployment gap.[25]

---

25. To our knowledge, there has been little empirical work estimating the role of the participation gap compared to the unemployment gap in affecting marginal costs and inflation. As noted earlier, our calibration seems consistent qualitatively with the implications of estimated insider–outsider models, although it seems most plausible that cyclical nonparticipants would put even less downward pressure on wages than the long-term unemployed.

We calibrate the adjustment cost parameter φ_G on labor force participation and the curvature parameter of the home production function σ_G to imply that the reduced-form parameters μ_1 and ϕ_m in the labor force participation equation (19) equal 0.92 and 0.07, respectively. These parameter estimates match our estimated coefficients for the post-1990 sample period that are reported in Table 3. The scaling parameters χ_0 and γ_G in the household utility function affecting the disutility of work and utility of home production, respectively, are set to imply a steady-state unemployment rate u of 5% (u = 0.05) and LFPR of $l_f = 2/3$. The corresponding values of the deep parameters that yield these implications are $\chi_0 = 1.37$, $\gamma_G = 0.039$, $\sigma_G = 2.82$, and $\phi_G = 1.46$.[26]

The price contract duration parameter ξ_p is set to 0.975, implying a very low degree of responsiveness of inflation to marginal cost of 0.0008, while the price indexation parameter ν is set to 0.8. We will later use the model to compare the performance of alternative policy rules following large adverse shocks, which cause policy rates to be pinned at the zero lower bound for a long duration. Given the simple structure of the model, long price contracts are required to allow the model to generate plausible variation in inflation, and also to solve the model numerically under some of the rules considered.[27] In future research, it would be useful to modify the model to include sticky wages as well as prices; wage rigidities would damp the responsiveness of inflation to unemployment, and allow for more realistic price contract durations.

The calibration of the monetary rule is described later. Table 4 provides a summary of parameter values under the benchmark calibration.

2.5 The Dynamics of the LFPR

We next illustrate how our model can account—at least qualitatively—for some of the stylized facts about labor force participation mentioned in the Introduction. In this vein, the upper panel of Figure 7 shows the effects of a fairly transient monetary policy shock ϵ_{it} in the policy rule (20) assuming that monetary policy is unconstrained by the zero lower bound. The monetary shock follows an AR(1) with a persistence of 0.7, and is scaled so that the policy rate rises 75 basis points at impact. The shock depresses employment by about 0.7% at impact, with habit persistence in consumption accounting for the slightly hump-shaped response.[28] Importantly, the

---

26. The Appendix describes the steady state of the model.

27. We solve the log-linearized model subject to the zero lower bound on policy rates using the numerical algorithm of Hebden, Linde, and Svensson (2013) and Erceg and Linde (2014). A solution is only regarded as admissible if it imposes the max operator in the interest rate reaction function. Our numerical analysis indicates that our ability to identify such solutions depends critically on the structure of the economy, including both the Phillips Curve slope and monetary policy rule. Even for calibrations that do not yield a solution satisfying the max operator, we can typically find a solution which implies that the policy rate remain at zero for say T periods (and then follows the usual law of motion implied by the Blanchard–Kahn conditions). The problem is that such solutions exhibit perverse dynamics similar to those described by Carlstrom, Fuerst, and Paustian (2012). In particular, contractionary fundamental shocks "that would depress output and inflation in a short-lived liquidity trap" can cause output and inflation to rise as the liquidity trap duration lengthens.

28. In Figure 7, the monetary rule is calibrated so that $\gamma_\pi = 2$ and $\gamma_u = 0.25$ (equal to unity if the interest rate is annualized). The coefficient on the unemployment gap may be interpreted as corresponding

TABLE 4

PARAMETERS UNDER BENCHMARK CALIBRATION

	Parameter	Value
Discount factor	β	0.995
Steady-state inflation	π	0.005
Intertemporal elasticity	σ_c	1.01
Taste shock scaling	ε_c	0.01
Capital share	α	0.30
Inverse Frisch elasticity	χ	4
Consumption habit	ψ_c	0.8
Participation gap persistence	μ_1	0.92
Participation gap sensitivity	ϕ_m	0.07
Price contract duration	ξ_p	0.975
Price indexation	ν	0.8
Inflation in policy rule	γ_π	2
Unemployment in policy rule	γ_u	0.25
Monetary shock persistence	ρ_i	0.7
Discount factor persistence	ρ_d	0.97
Consumption taste shock persistence	ρ_c	0.85

NOTES: The full stabilization Taylor rules considered in Figures 8–13 impose an arbitrarily high coefficient on the variable stabilized by the rule, so that, for example, $\gamma_u = 150$ in the case in which the unemployment gap is fully stabilized. The following alternative parameter values are also used in the figures:
Figure 7: $\rho_i = 0.9$ under the "persistent tightening" case.
Figure 9: $\mu_1 = 0.8$ and $\phi_m = 0.186$ under both of the monetary rules considered.
Figure 10: Monetary rule under "commitment" case based on loss function described in Section 3.3.
Figure 11: $\xi_p = 0.94$ under the "faster price adjustment" case.
Figure 12: $\chi = 1$ and $\xi_p = 0.96$ under the "large response" case; and $\chi = 100$ and $\xi_p = 0.996$ under the "tiny response" case.

substantial near-term fall in employment n_t is driven almost entirely by a rise in the unemployment rate (not shown, but is given by the gap between employment and labor force participation given that our simulations report the scaled unemployment gap $\frac{1}{1-u}(u_t - u_t^*)$). Labor force participation does exhibit a very persistent longer run decline that contributes to a "trendlike" decline in the employment gap, but this variation is small compared to the near-term effect.

Given that labor force participation simply follows an AR(1) in our benchmark specification—with the unemployment rate the forcing variable:

$$l_{ft} = \mu_1 l_{ft-1} - \phi_m u_t, \tag{25}$$

labor force participation can exhibit a large eventual drag on employment if shocks keep unemployment persistently above baseline. This is illustrated in the lower panel for a persistent monetary tightening implemented through a sequence of monetary policy innovations lasting for 3 years (with the shocks having a persistence of 0.9). About half of the employment gap at a horizon of 5 years is attributable to a drop in labor force participation, which seems roughly in line with the U.S. experience in the aftermath of the Great Recession. Moreover, as the economy recovers and the unemployment rate starts to fall (note that the gap between the lines begins narrowing

to the standard Taylor rule coefficient of 0.5 on the output gap under the assumption of an Okun's law coefficient of 0.5 relating the unemployment gap to the output gap.
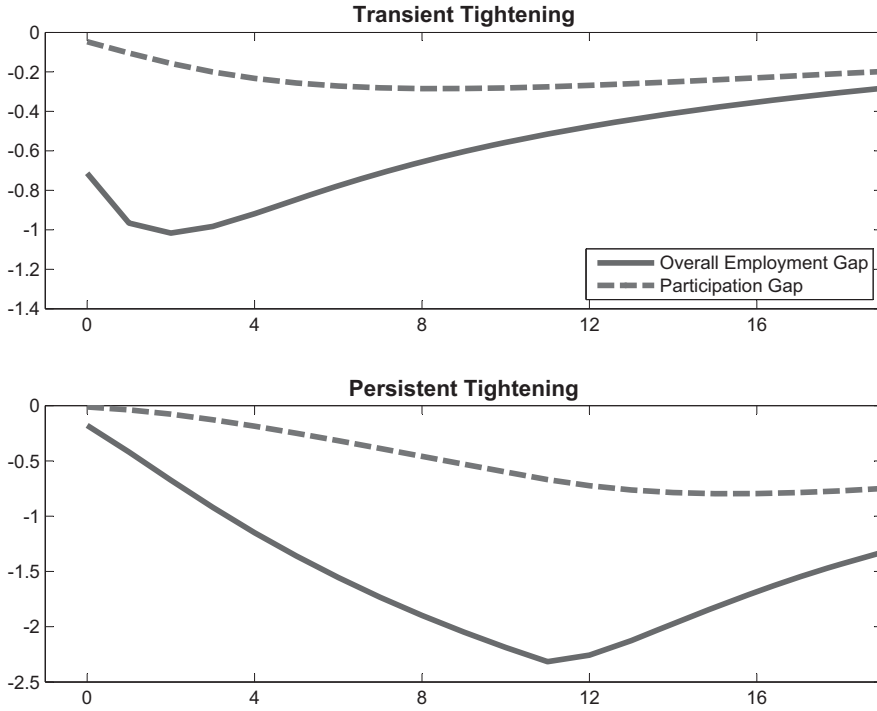
Transient Tightening

Persistent Tightening

Fɪɢ. 7. The Response of the Labor Market to Monetary Policy Shocks.

after 12 quarters), labor force participation continues to decline and accounts for progressively more of the employment gap. We consider later how this divergence between gaps can have major implications for monetary policy.

As is clear from equation (25), real shocks that reduced demand very persistently would induce a similar path for labor force participation. Even so, the case of a protracted monetary tightening is insightful insofar as it captures how a protracted period at the zero lower bound—which operates like a sequence of adverse monetary shocks—can contribute to what appears to be a trendlike deterioration in employment.

3. THE DESIGN OF MONETARY POLICY RULES

We begin this section with an overview of how monetary policy trade-offs in our model are influenced by the LFPR; this analysis is quite general insofar as it does not hinge on our particular calibration of key parameters.

3.1 Policy Trade-Offs

A familiar result based on the workhorse New Keynesian model—highlighted by Goodfriend and King (1997) and dubbed "divine coincidence" by Blanchard and Gali (2010)—is that monetary policy faces no trade-off between stabilizing inflation and unemployment. Thus, in response to an adverse demand shock that temporarily boosted the unemployment rate above its natural rate, inflation would return to target as soon as u_t reverted to u_t^*, and a policy of strictly targeting the unemployment gap would yield the same outcome as a policy that strictly targeted inflation.[29] As we discuss later, the workhorse model may be interpreted as making the implicit assumption that the participation gap moves *contemporaneously* with unemployment, so that the "no trade-off" result also applies to the employment gap.

In our framework, a large negative labor force participation gap—presumably caused by a deep recession—can markedly affect monetary policy trade-offs relative to the workhorse New Keynesian model, at least under the plausible conditions satisfied by our benchmark calibration. One condition is that labor force participation responds gradually to the unemployment rate: recalling Figure 7, a sluggish response means that labor force participation can remain depressed even as the economy recovers and unemployment rate approaches u_t^*. A second condition is that the unemployment gap affects inflation by relatively more than the participation gap, so $\psi_u > \psi_l$.

Assuming these conditions hold, there are several key dimensions along which the implications of our model diverge from the workhorse New Keynesian model. First, as can be inferred from equation (17), stabilizing the unemployment gap ($u_t = u_t^*$) does not suffice in stabilizing inflation at target if the participation gap is negative (or more generally, nonzero). This reflects that a negative labor force participation gap puts at least some downward pressure on marginal cost and hence inflation even if $u_t = u_t^*$.

A second and closely related implication of our model is that a policy of aggressively targeting inflation in the wake of a persistently negative labor force participation gap requires the unemployment rate to "overshoot," that is, fall below the long-run natural rate for some time. To see this, note that the Phillips Curve equation (16) requires that discounted marginal cost must be zero ($\sum_{j=0}^{\infty} \beta^j mc_{t+j} = 0$) for inflation to be stable, which implies:

$$\sum_{j=0}^{\infty} \beta^j \psi_u \frac{1}{1-u} (u_{t+j} - u_{t+j}^*) = \sum_{j=0}^{\infty} \beta^j \psi_l (l_{ft+j} - l_{ft+j}^*), \qquad (26)$$

where for simplicity the foregoing expression assumes no structural persistence in inflation. Thus, the unemployment rate must on average overshoot the natural rate ($u_t < u_t^*$) by enough to balance the deflationary pressure arising from current and

---

29. We are assuming implicitly that monetary policy is forced to depart from the strict targeting rule for some time—say, due to a binding zero bound constraint—so that the demand shock boosts unemployment.

expected future negative participation gaps (the right-hand side of equation (26)). The implication that unemployment must fall below the natural rate to keep inflation stable in our model is tantamount to a fall in the short-term natural rate of unemployment.[30]

A third key implication is that a policy of aggressively targeting the employment gap puts upward pressure on inflation. From equation (15), setting $n_t = n_t^*$ requires that:

$$\frac{1}{1-u}(u_t - u_t^*) = -(l_{ft} - l_{ft}^*).$$ (27)

In other words, any shortfall in participation ($l_{ft} < l_{ft}^*$) must be fully offset by a fall in unemployment relative to the natural rate. To the extent that the negative participation gap puts less downward pressure on inflation than the unemployment gap, so $\psi_l < \psi_u$ (assuming both gaps are scaled to have similar effects on the employment gap), balancing the gaps equally as implied by (27) puts upward pressure on inflation. Thus, as we explore further later, while monetary policy can achieve a faster recovery by targeting the employment gap, such a policy typically involves the trade-off of higher inflation.

Our model implies that there is no trade-off between employment gap and inflation stabilization in three special cases; while somewhat extreme, they are nonetheless useful for gauging how monetary policy trade-offs are likely to vary as the calibration is modified along these dimensions. One case occurs if labor force participation responds contemporaneously to the current unemployment gap, which is implied if the adjustment cost parameter $\varphi_G = 0$. Because the unemployment gap and participation gaps are each simply proportional to the employment gap $n_t - n_t^*$, marginal cost and hence inflation also depend only on $n_t - n_t^*$, and our model becomes observationally equivalent to the New Keynesian model. Accordingly, monetary policy would face no trade-off between stabilizing inflation and the employment gap, and it would be immaterial whether monetary policy responded to the unemployment gap or employment gap. Quite intuitively, this special case highlights that the choice of a gap to target—the unemployment gap versus the employment gap—should be less consequential to the extent that the participation gap responds relatively quickly to unemployment.

Second, even assuming positive adjustment costs, there would be no trade-off between employment gap and inflation stabilization if the participation gap was initially zero, as would occur if monetary policy had always kept unemployment at its natural rate in the past. This highlights how the zero lower bound—through causing a participation gap to emerge—may give rise to policy trade-offs.

Finally, there is no trade-off between stabilizing the employment gap and inflation if $\psi_l = \psi_u$.[31] In this case, the upward pressure on inflation associated with pushing unemployment below the natural rate is exactly offset by downward pressure

---

30. Clearly, this contrasts with the "no trade-off" result in the New Keynesian model, which implies that inflation is stabilized through bringing unemployment back to its natural rate (with no required overshooting).

31. In our model, this occurs as the Frisch elasticity of labor supply tends toward infinity.

arising from the cyclical nonparticipants.[32] As suggested by this special case (and demonstrated later), a relatively large coefficient on the participation gap (ψ_l close to ψ_u) means that employment gap targeting should be associated with less inflationary pressure.

3.2 Simple Policy Rules

Our qualitative analysis mentioned above suggests that monetary policy must allow the unemployment rate to fall below its long-run natural rate to keep inflation stable, but faces a potential trade-off in pursuing a more aggressive policy of stabilizing the employment gap. We next illustrate some dynamic implications using simulations of our model under a noninertial Taylor rule. We generate initial conditions for our simulations by assuming that the economy is buffeted by a persistent adverse demand shock that pins the policy rate i_t at its zero lower bound for roughly 2 years. In all of our simulations, the shock consists of a persistent change in the discount factor ε_{dt} that causes the natural real interest rate r_t^* to fall sharply.

The dash-dotted lines in Figure 8 show the effects of the demand shock under a Taylor rule that aims to fully stabilize the unemployment rate at its long-run potential rate ($u_t = u_t^*$); thus, the Taylor rule given by equation (20) is calibrated with a very large coefficient γ_u that would completely insulate the real economy from this demand shock if the zero bound constraint were not binding (we set $\gamma_u = 150$). Because the zero bound does bind in our simulation, the unemployment rate u_t rises far above u_t^* (panel 1), the participation rate falls persistently (panel 2), and inflation declines well below target (panel 4). As the shock wears off and the natural real interest rate rises, this policy rule implies a smooth convergence of u_t to u_t^*.[33] Moreover, while the unemployment rate gap converges within 10 quarters, the effects on labor force participation are much longer lived. The highly persistent negative participation gap puts downward pressure on inflation, so that inflation remains below target well after the unemployment rate gap has closed.

Figure 8 also shows the implications of alternative policy rules aimed at strictly stabilizing inflation—the dashed lines—or the overall employment gap $n_t - n_t^*$—the solid lines (achieved by setting γ_π or γ_n to arbitrarily large values, respectively). Consistent with our previous discussion, the rule that attempts to stabilize inflation implies an eventual undershooting of the unemployment rate gap as the economy recovers (i.e., $u_t < u_t^*$). Quite intuitively, pushing the unemployment rate below the long-run NAIRU helps counterbalance the downward pressure on inflation arising from a persistently negative participation gap. The rule that attempts to stabilize the overall employment gap n_t is even more aggressive insofar as it aims to allow the

---

32. This special case highlights how the "no trade-off" result applies to the employment gap, and not to the unemployment gap (as in the workhorse New Keynesian model). A large coefficient ψ_l on the participation gap in the marginal cost relation would imply substantial downward pressure on inflation under a rule targeting the unemployment gap.

33. The figures plots interest rates and inflation in levels, mainly to highlight the zero bound constraint on nominal rates. Real variables are simply depicted as a percent deviation from their steady-state level.
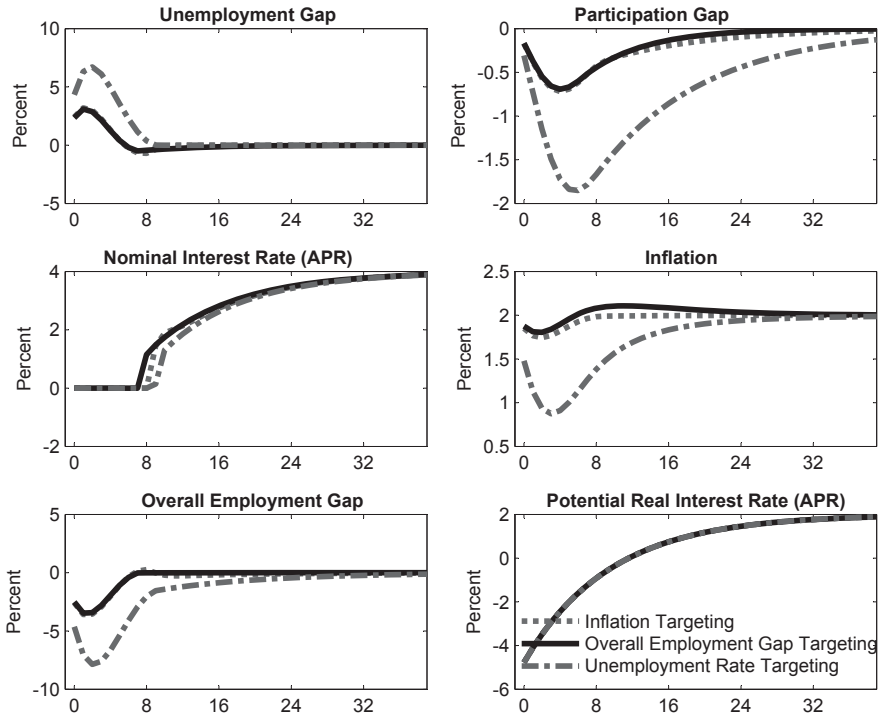
Fɪɢ. 8. Model Behavior under Alternative Policy Rules.

unemployment rate to overshoot by enough to *fully offset* the negative participation gap (recalling equation (15)). Given the larger coefficient on the unemployment gap relative to the participation gap in the Phillips Curve, inflation rises persistently above target as the economy recovers. This policy rule not only promotes a faster recovery in both unemployment and participation than the other rules considered, but also mitigates the initial downturn in employment and inflation following the adverse shock.

While the calibration underlying Figure 8 implies a large negative negative participation gap even after the unemployment rate has recovered—an implication that seems consistent with the Great Recession—there is admittedly considerable uncertainty about how labor force participation would respond to a recovering economy. To the extent that participation responds more quickly to unemployment (μ_1 is lower), the disparity between a rule that responds to employment versus unemployment becomes smaller.[34] This is illustrated in Figure 9, which assumes that the persistence parameter μ_1 is set at a substantially lower value of 0.8; the calibration also sets $\phi_m = 0.19$ to

---

34. As noted, the difference disappears completely in the special (though unrealistic) case of no adjustment costs.
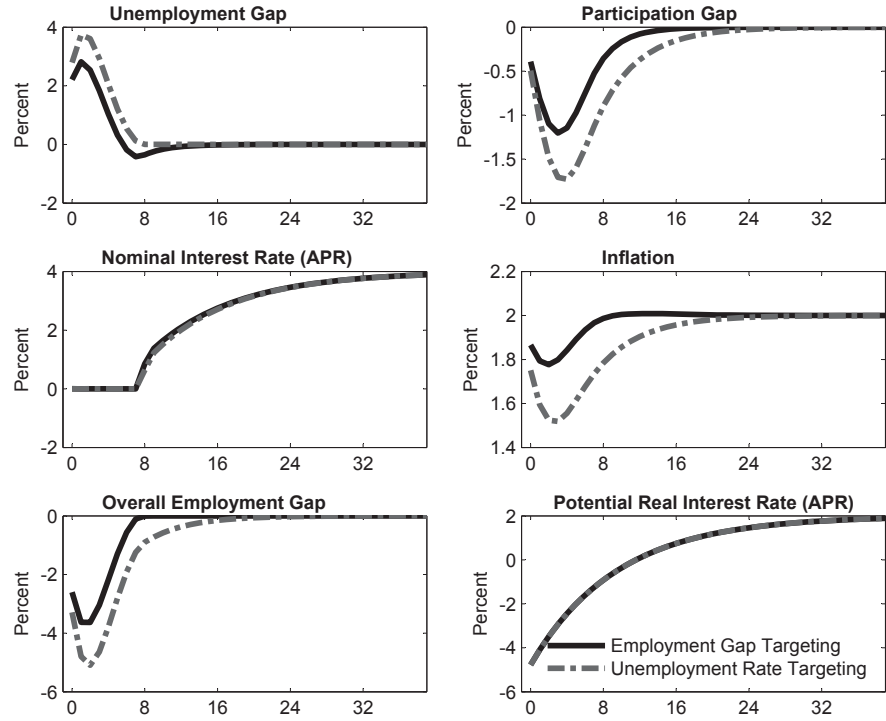
Fig. 9. Implications of a Less Inertial Participation Gap.

keep the long-run response of participation to unemployment unchanged from the benchmark. Under this alternative calibration, the participation rate rebounds fairly quickly even under unemployment gap targeting. Conversely, the comparative advantage of employment gap targeting in ameliorating the recession would be even larger than that shown in (the previous) Figure 8 if labor force participation responded more sluggishly than under our benchmark calibration; of course, a potential cost would be a greater overshooting of inflation.

The disparity between the speed of the recovery under employment gap targeting and unemployment targeting would presumably widen relative to our benchmark to the extent that our model was augmented to allow for relevant endogenous propagation channels. For example, in a model with financial frictions, the faster recovery under employment gap targeting would be accompanied by a somewhat quicker decline in credit spreads, and similar results would obtain in a modeling environment that allowed for precautionary savings.

The implication that monetary policy can spur a faster recovery in slow-moving state variables by responding to the gap between a stock variable and its target level extends well beyond our model in which the participation rate is the only endogenous

state variable. For example, in a model with endogenous capital accumulation, the capital stock would recover more slowly than control variables such as investment or hours worked (much like the LFPR in the lower panel of Figure 7). A monetary policy rule that in effect responded to the "capital stock gap"—the deviation of the capital stock from its level under flexible prices—would induce the same sort of overshooting of controls as evident in Figure 8 for the unemployment rate, and have similar qualitative implications for inflation. Characterizing such trade-offs in an empirically realistic model would seem a fruitful subject for future research. Of course, the extent to which a capital stock gap or housing gap influenced the course of policy would depend heavily on policymaker objectives. Thus, a central bank's willingness to tolerate above-target inflation to close an employment gap may differ considerably from its willingness to accept higher inflation in order to close a capital stock gap, particularly if the central bank viewed its objectives as more closely linked to employment.

3.3 Commitment-Based Strategies

The implications of the employment targeting rule shown in Figure 8—in which unemployment eventually falls below its natural rate and inflation rises above target— are broadly similar to those implied by the optimal commitment policy in the face of the zero lower bound (cf. Eggertsson and Woodford 2003). Our results show that a noninertial policy rule may induce undershooting of unemployment along with some overshooting of inflation. In effect, a policymaker in our model who cares about inflation and total employment—not just the unemployment rate—would be inclined to follow a policy that allows inflation to exceed its target for a while in order to promote a faster narrowing of the participation gap.

Of course, the employment gap targeting rule considered above is noninertial and hence cannot capture potentially substantial gains that may be associated with a commitment-based strategy. While deriving a utility-based welfare function is beyond the scope of this paper, the potential benefits of a commitment-based strategy can be illustrated with a quadratic loss function that simply involves the squared employment gap and square inflation gap.[35] We then proceed to derive the targeting rule under commitment that optimizes this objective function; the implied policy rule embeds history dependence (as in the literature on commitment-based strategies) and hence provides a contrast with the noninertial rules considered earlier.

As seen in Figure 10, the commitment-based policy rule involves an eventual overshooting of the employment gap and thereby provides more stimulus than a noninertial rule. Consequently, the unemployment rate falls even further below its natural rate, while the participation rate exhibits overshooting as the economy recovers. This behavior clearly contrasts with the noninertial employment targeting

---

35. The weights in the loss function on inflation and employment are assumed to be equal if that inflation is expressed at an annualized rate; for technical reasons, we also place a small weight on the change in the short-term nominal interest rate.
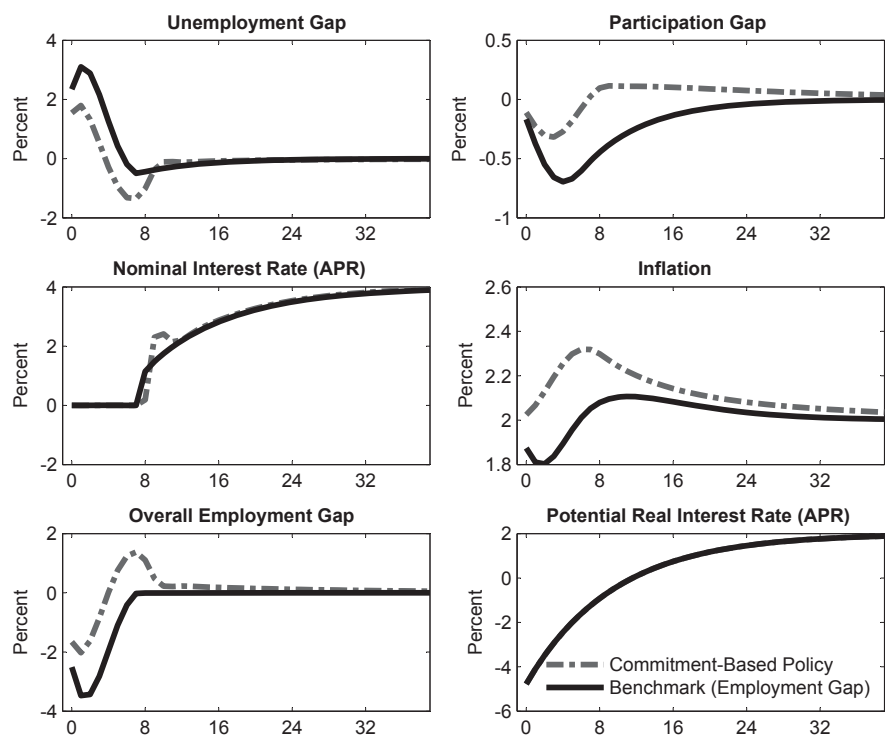
FIG. 10. Policy under Commitment with an Ad Hoc Loss Function.

rule, which implies monotonic convergence of the employment gap and participation gap, albeit at different speeds.

3.4 Uncertainty and Risks

Under the benchmark calibration shown in Figure 8, employment gap targeting generates a relatively fast recovery in the labor market, while pushing inflation only moderately above target. Even so, policymakers might be concerned about the possibility of less benign outcomes for inflation if the structure of the economy turned out to be different than under our benchmark calibration or, alternatively, if the participation gap could only be estimated imprecisely in real time. Hence, we next consider some key factors affecting the inflation response under employment gap targeting in environments that depart from our benchmark model. Our discussion should be interpreted as more broadly applicable to policies, which put some weight on the participation gap.

One key factor affecting the inflation response is the slope of the Phillips Curve. Inflation could rise by substantially more as the economy recovered if the Phillips
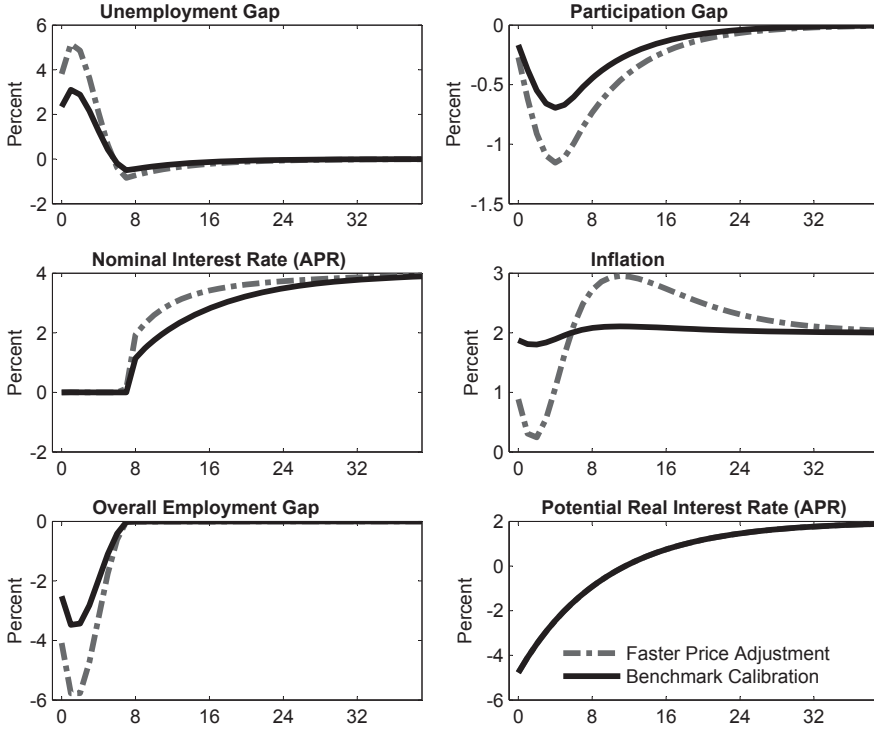
Fɪɢ. 11. Implications of a Steeper New Keynesian Phillips Curve.

Curve slope were steeper. This scenario is illustrated in Figure 11, which compares our benchmark calibration to an alternative in which the Phillips Curve slope parameter κ_p is several times larger.[36] Inflation rises to around 3% under this alternative and remains persistently above target.

As shown in Figure 12, the response of inflation under the employment gap targeting rule also depends on the degree to which wages respond to the participation gap. First, we consider the case where individuals who have dropped out of the labor force have no influence at all on current inflation; that is, we set the NKPC coefficient on the participation gap ψ_l arbitrarily close to zero (by choosing a high value of χ, the inverse Frisch elasticity). Evidently, the trajectory for inflation is only slightly higher in this scenario, reflecting that this coefficient has a relatively small value even under our benchmark calibration. Next, we consider the case in which wages are much more sensitive to the participation gap. In particular, we set the Frisch elasticity of labor supply ($\frac{1}{\chi}$) equal to unity, which implies that wages are about half as responsive

36. This is achieved by setting the contract duration parameter $\xi_p = 0.94$ rather than 0.975 as in our baseline.
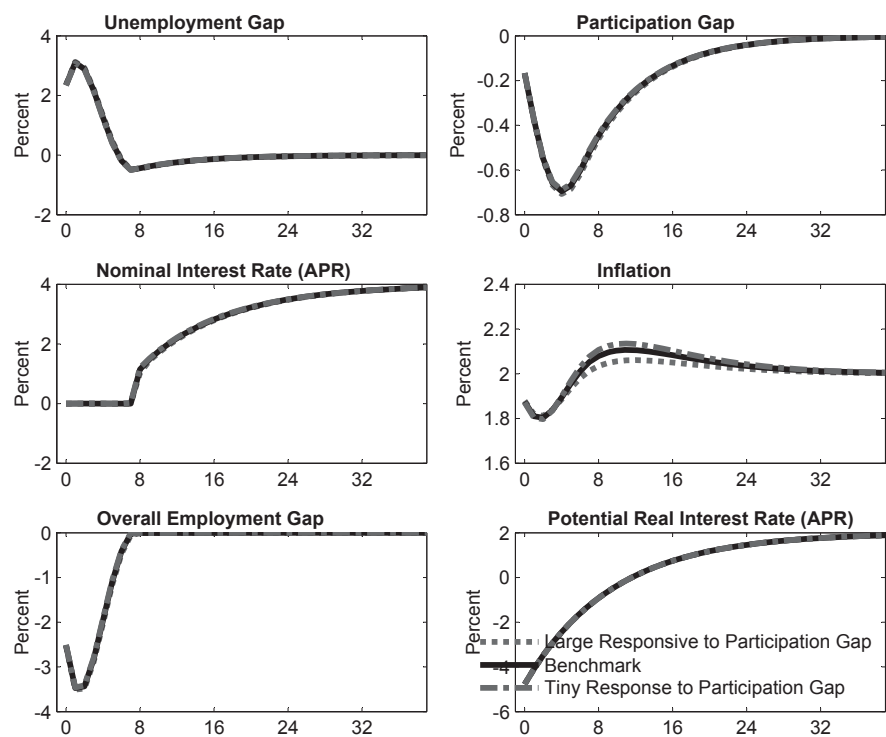
FIG. 12. Inflation Responsiveness to the Participation Gap.

to the participation gap as to the unemployment gap. In this case, the employment gap targeting rule keeps inflation very close to target; that is, there is effectively little trade-off in promoting these objectives.

The inflation trade-offs associated with a strategy of employment gap targeting also depend on the policymaker's ability to estimate the participation gap. In our model, both the private sector and central bank correctly infer l_{ft}^*, the participation rate that would prevail in an environment in which prices were always flexible, and the central bank's policy rule aims to bring participation (and unemployment) back to its natural rate.[37] However, policymakers in reality face considerable uncertainty in decomposing labor force participation into "cyclical" and "trend" components, just as they face challenges in decomposing unemployment fluctuations. In particular, while we think our empirical analysis in Section 1 presents a strong case that much of the post-2007 fall in the LFPR is cyclical, there is a range of plausible estimates of the *current* participation gap. Moreover, there is also considerable uncertainty about how the participation gap will evolve as the economy recovers, which can be

---

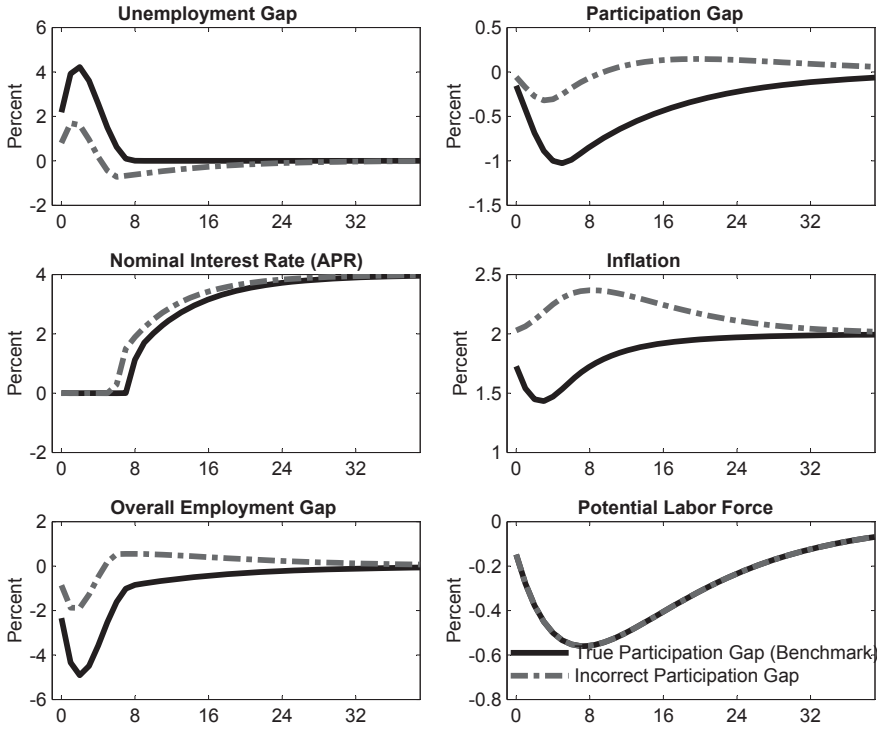37. This interpretation of the natural rate of unemployment and participation follows Neiss and Nelson (2003).

FIG. 13. Policy Responds to Alternative Participation Gap Measures.

regarded as uncertainty about the parameters μ_1 and ϕ_m in the context of our stylized model; as we have discussed, a more gradual response of participation causes more inflation overshooting under employment gap targeting. Finally, policymakers face the additional challenge of having to decompose participation into cycle and trend in real time.

In line with an extensive literature emphasizing how high inflation can result if policymakers attempt to target an unemployment rate below the true NAIRU—including seminal work by Orphanides (2002)—the risk of generating substantial inflationary pressure under a strategy that responds to the participation gap would seem highest if policymakers consistently overestimated the natural (or potential) participation rate l_{ft}^*. In the case of the recession scenarios considered, this would be tantamount to overestimating the cyclical component of the participation gap (in absolute terms), and hence underestimating the decline due to secular influences (which could include shocks that would affect the desirability of working even in the absence of sticky prices or wages).

To illustrate potential inflation risks, Figure 13 shows a scenario in which a real shock—a fall in the marginal utility of consumption—reduces the return to working

in the market sector, and thus lowers both the natural participation rate l_{ft}^* and the natural employment rate (i.e., raises u_t^*).[38] In this setting, a monetary policy rule that aims to bring the LFPR back to its preshock level induces considerably more upward pressure on inflation than under the alternative in which monetary policy responds to the "true" participation gap $l_{ft} - l_{ft}^*$.

There are various options for addressing inflation risks associated with a strategy that involves responding to the participation gap, some of which parallel the literature on natural rate misperceptions. For example, policymakers may adopt a very conservative measure of the cyclical component of the participation gap that enters into the monetary policy rule or, alternatively, may simply reduce the weight on the participation gap in the monetary rule in a way proportional to measurement uncertainty. However, these approaches have the obvious drawback that they may contribute to a much slower labor market recovery. Given these costs, it would seem fruitful to devote much more attention in future research to attempting to measure the cyclical component of the participation gap, perhaps through focusing on somewhat restricted segments of the population for which for which demographic shifts are less pronounced (e.g., prime-aged workers). While there are indeed significant challenges in measuring the cyclical component of participation, it is not clear that these challenges are any more formidable than those associated with measuring the natural rate of unemployment.

4. CONCLUSIONS

In this paper, we have drawn on both U.S. and state-level evidence to show that much of the decline in the LFPR since 2007 is due to cyclical factors, and then used a simple extension of the workhorse New Keynesian model to investigate the monetary policy trade-offs associated with responding to broader measures of resource slack that include participation. Monetary policy can induce a faster narrowing of the employment gap by allowing the unemployment rate to significantly fall below the natural rate; however, pushing unemployment below the natural rate may put upward pressure on wages and hence inflation, so that there is a potential trade-off.

We have used our model to highlight key factors influencing this trade-off, including the slope of the Phillips Curve, the sensitivity of market wages to the participation gap, and the dynamic response of participation to unemployment. However, relatively little empirical work has been done to characterize how market wages are affected by cyclical movements in participation, or how participation responds to long periods of labor market slack. Although this paucity of evidence is unsurprising given limited experience with deep recessions, more empirical work in these areas clearly seems

---

38. Our scenario assumes that the central bank responds to the "true" natural rate u_t^*—which rises in response to this shock—in order to highlight the consequences of misestimating the participation gap alone; clearly, the effects on inflation would be even larger if the central bank assumed that u_t^* was unaffected.

warranted. In this vein, microdata—including time use surveys—may be invaluable in helping illuminate the factors that influence participation decisions.

There are also many factors likely to affect monetary policy trade-offs that are not captured by our simple modeling framework but that seem highly relevant in practice for crafting an appropriate strategy. Given the very stylized nature of our model, it would be useful to introduce some of the key features incorporated into more empirically realistic DSGE models, including nominal wage rigidities, hand-to-mouth spending behavior, endogenous capital accumulation, and financial frictions. As noted, the differences between alternative monetary policy rules would likely be accentuated by such endogenous propagation channels. It would also be desirable to reconsider our analysis in a framework that includes explicit labor market search and that takes account of how features such as skill deterioration associated with long-term unemployment influence the job matching process.

APPENDIX

The steady-state values of output Y, the employment of the population N, the unemployment rate of the labor force u, the LFPR L_f, and the real wage $\frac{W}{P}$ are determined by the following equations:

$$Y = N^{1-\alpha}, \tag{A1}$$

$$N = H_m L_f = (1 - u)L_f, \tag{A2}$$

$$(1 - \alpha)(L_f(1 - u))^{-\alpha} = \frac{W}{P}, \tag{A3}$$

$$\chi_0(1 - u)^{\chi} = \frac{W}{P}((1 - \psi_c - \varepsilon_c)y)^{-\sigma_c}, \tag{A4}$$

$$\chi_0(1 - u)^{1+\chi} \frac{\chi}{1 + \chi} = \gamma_G(1 - L_f)^{-\sigma_G}. \tag{A5}$$

The first equation is simply the aggregate production function (assuming the capital stock is unity). The second equation expresses the employment rate of the population as the product of the employment rate of the labor force (i.e., $1 - u = H_m$) and the LFPR. The third equation is the firm's labor demand schedule, and the fourth equation the labor supply schedule for a household in the market sector (where the marginal utility of consumption equals $((1 - \psi_c - \varepsilon_c)y)^{-\sigma_c}$, recalling that $C = Y$). The fifth condition determines the allocation of households between the market and

home production sectors. Finally, we note that the composite parameter μ_1 in the evolution equation for the LFPR (25) can be expressed as:

$$\mu_1 = \frac{\hat{\phi}_G}{1 + \hat{\phi}_G}, \tag{A6}$$

where $\hat{\phi}_G$ itself depends on the adjustment cost parameter ϕ_G according to:

$$\hat{\phi}_G = \frac{\phi_G}{\gamma_G \sigma_G} \frac{(1 - L_f)^{1+\sigma_G}}{L_f}((1 - \psi_c - \varepsilon_c)y)^{\sigma_c}. \tag{A7}$$
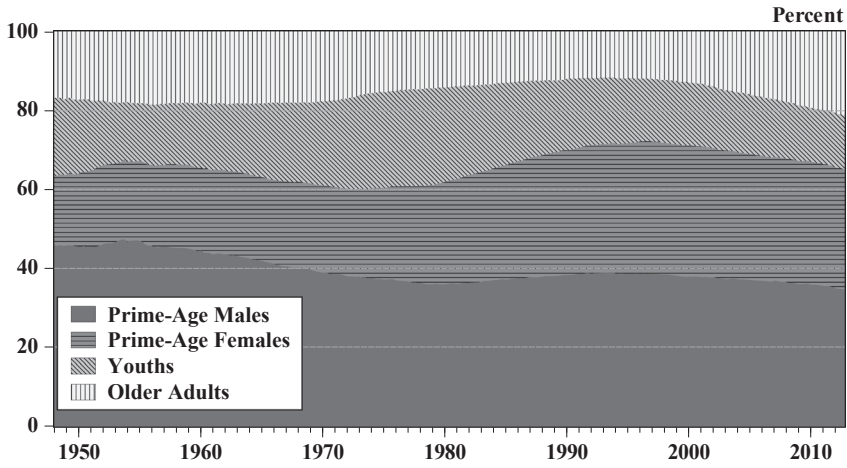

Fig. A1. Demographic Composition of the Labor Force.

NOTES: This figure depicts the demographic composition of the labor force (in percent) from 1948 to 2012. The solid area denotes the share of prime-age males (25–54 years old), and the horizontal-dashed area denotes the share of prime-aged females (25–54 years). The diagonal-dashed and vertical-dashed areas denote the shares for youths (16–24 years) and older adults (55 years and above), respectively.

Males, 16 to 24 Years

Females, 16 to 24 Years

BLS forecast (Nov2007)

Males, 25 to 54 Years

Females, 25 to 54 Years

Males, 55 Years and Older

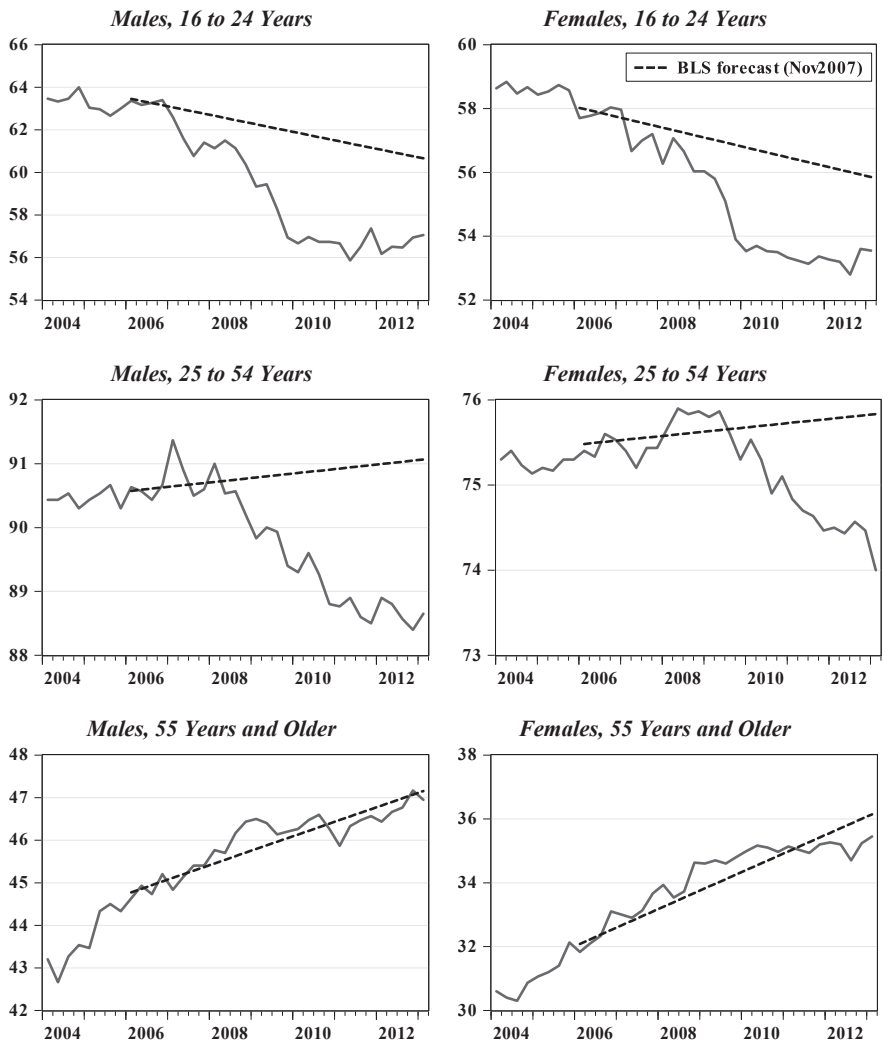### Females, 55 Years and Older

FIG. A2. Labor Force Participation by Age and Gender.

NOTES: In each panel, the solid line denotes the actual evolution of the labor force participation rate (in percent) for the specified demographic group from 2004:Q1 to 2013:Q1, and the dashed line denotes the linear trend for the decade ending in 2016 that was projected by BLS staff and published in the November 2007 issue of the *Monthly Labor Review*.
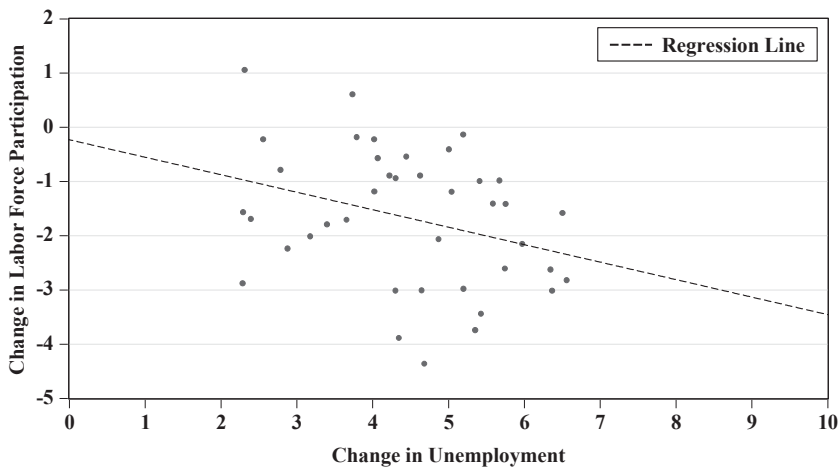
FIG. A3. Data from Larger States (Excluding Nevada) Regarding Unemployment and Labor Force Participation of Prime-Age Adults.

NOTES: The vertical axis refers to the change in the labor force participation rate for prime-age adults (25–54 years) between 2007 and 2012, and the horizontal axis refers to the change in the unemployment rate for that demographic group between 2007 and 2010. The observations correspond to 40 states that had more than 500,000 prime-age adults as of December 2007; that is, this sample excludes Nevada (for reasons discussed in the text) as well as Alaska, Delaware, Hawaii, Montana, North Dakota, Rhode Island, South Dakota, Vermont, Wyoming, and Washington, DC. The dashed line depicts the regression obtained using these 40 data points.
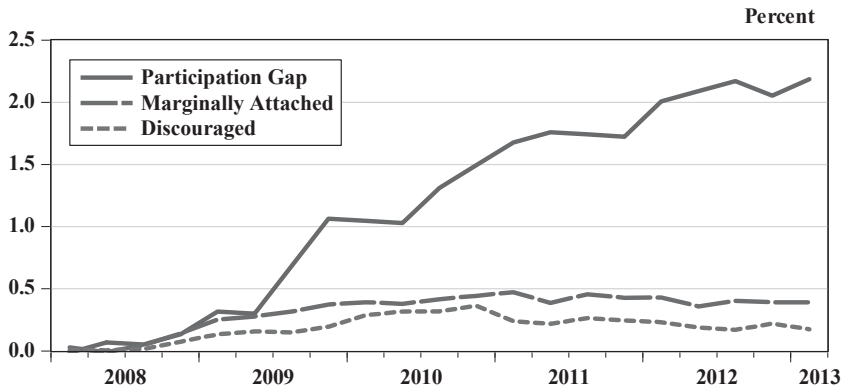

FIG. A4. Gauging the Degree of Attachment to the Labor Force.

NOTES: The solid line denotes the deviation between the actual LFPR and its trend path as projected by the BLS in November 2007. The short-dashed and long-dashed lines denote the post-2007 changes in the incidence of "discouraged job-seekers" and in the incidence of individuals who are "marginally attached to the labor force," respectively, measured as a share of the civilian noninstitutionalized population aged 16 and above. "Marginally attached" refers to persons not in the labor force who want a job, are available now, and have looked for work sometime in the past year but not in the past month. "Discouraged" refers to persons not in the labor force who want a job and are available now and who have looked for work sometime in the past year but who are not currently searching because they believe there are no jobs available for which they would qualify.

LITERATURE CITED

Aaronson, Daniel, Jonathan Davis, and Luojia Hu. (2012) "Explaining the Decline in the U.S. Labor Force Participation Rate." Chicago Fed Letter 296, Federal Reserve Bank of Chicago.

Aaronson, Daniel, Kyung Park, and Daniel Sullivan. (2006) "The Decline in Teen Labor Force Participation." *Economic Perspectives, Federal Reserve Bank of Chicago*, 30, 2–18.

Aaronson, Stephanie, Bruce Fallick, Andrew Figure, Jonathan Pingle, and William Wascher. (2006) "The Recent Decline in the Labor Force Participation Rate and Its Implications for Potential Labor Supply." *Brookings Papers on Economic Activity*, 69–134.

Autor, David. (2010) *The Polarization of Job Opportunities in the U.S. Labor Market*. Washington, DC: Center for American Progress.

Autor, David, and Mark Duggan. (2003) "The Rise in the Disability Rolls and the Decline in Unemployment." *Quarterly Journal of Economics*, 118, 157–205.

Autor, David, and Mark Duggan. (2006) "The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding." *Journal of Economic Perspectives*, 20, 71–96.

Benhabib, Jess, Richard Rogerson, and Randall Wright. (1991) "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." *Journal of Political Economy*, 99, 1166–87.

Blanchard, Olivier. (1991) "Wage Bargaining and Unemployment Persistence." *Journal of Money, Credit, and Banking*, 23(Part 1), 277–92.

Blanchard, Olivier, and Peter Diamond. (1994) "Ranking, Unemployment Duration, and Wages." *Review of Economic Studies*, 61, 417–34.

Blanchard, Olivier, and Jordi Gali. (2010) "Labor Markets and Monetary Policy: A New-Keynesian Model with Unemployment." *American Economic Journal: Macroeconomics*, 2, 1–30.

Blanchard, Olivier, and Lawrence Summers. (1986) "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual 1986*, Vol. 1, edited by S. Fischer, pp. 15–78. Cambridge, MA: MIT Press.

Bound, John, and Timothy Waidmann. (2002) "Accounting for Recent Declines in Employment Rates among Working-Aged Men and Women with Disabilities." *Journal of Human Resources*, 37, 231–50.

Carlstrom, Charles, Timothy Fuerst, and Matthias Paustian. (2012) "Inflation and Output in New Keynesian Models with a Transient Interest Rate Peg." Bank of England Working Paper No. 459.

Chang, Yongsung, Joao Gomes, and Francis Schorfheide. (2002) "Learning-by-Doing as a Propagation Mechanism." *American Economic Review*, 92, 1498–1520.

Clark, Kim, and Lawrence Summers. (1982) "Labor Force Participation: Timing and Persistence." *Review of Economic Studies*, 49, 825–44.

Congressional Budget Office. (2007) "The Budget and Economic Outlook: Fiscal Years 2008-2017." January.

Congressional Budget Office. (2013) "The Budget and Economic Outlook: Fiscal Years 2013-2022." February.

Daly, Mary, Early Elias, Bart Hobijn, and Oscar Jorda. (2012) "Will the Jobless Rate Drop Take a Break?" Economic Letter 2012-37, Federal Reserve Bank of San Francisco.

DiCecio, Riccardo, Kristie Engemann, Michael Owyang, and Christopher Wheeler. (2008) "Changing Trends in the Labor Force: A Survey." *Federal Reserve Bank of St. Louis, Review*, January/February, 47–62.

Eggertsson, Gauti, and Michael Woodford. (2003) "The Zero Interest-Rate Bound and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, 1, 139–211.

Erceg, Christopher and Jesper Linde. (2014) "Is There a Fiscal Free Lunch in a Liquidity Trap?" *Journal of the European Economic Association*, 12, 73–107.

Federal Reserve Board. (2007) "FOMC Material for Staff Presentation on the Economic Outlook." January 30.

Gali, Jordi. (2011) "The Return of the Wage Phillips Curve." *Journal of the European Economic Association*, 9, 436–61.

Goldin, Claudia. (2006) "The Quiet Revolution that Transformed Women's Employment, Education, and Family." *American Economic Review*, 96, 1–21.

Goodfriend, Marvin, and Robert King. (1997) "The New Neoclassical Synthesis and the Role of Monetary Policy." In *The NBER Macroannual 1997*, edited by B. Bernanke and J. Rotemberg, pp. 231–82. Cambridge, MA: MIT Press.

Greenwood, Jeremy, and Zvi Hercowit. (1991) "The Allocation of Capital and Time over the Business Cycle." *Journal of Political Economy*, 99, 1188–1214.

Gruber, Jonathan. (2000) "Disability Insurance Benefits and Labor Supply." *Journal of Political Economy*, 58, 1162–83.

Hebden, James, Jesper Linde, and Lars Svensson. (2013) "Optimal Monetary Policy in the Hybrid New Keynesian Model under the Zero Bound Constraint." Manuscript, Federal Reserve Board.

Heckman, James, and Mark Killingsworth. (1986) "Female Labor Supply: A Survey." In *Handbook of Labor Economics*, Vol. 1, edited by O. Ashenfelter and R. Layard, pp. 103–204. Amsterdam: Elsevier Science Publishers.

Heim, Bradley. (2007) "The Incredible Shrinking Elasticities: Married Female Labor Supply 1978-2002." *Journal of Human Resources*, 42, 881–918.

Hotchkiss, Julie. (2005) "Changes in Behavioral and Characteristic Determination of Female Labor Force Participation, 1975-2005." *Economic Review*, *Federal Reserve Bank of Atlanta*, 1–13.

Hotchkiss, Julie, Melinda Pitts, and Fernando Rios-Avila. (2012) "A Closer Look at Non-Participants during and after the Great Recession." Federal Reserve Bank of Atlanta Working Paper No. 2012-10.

Hotchkiss, Julie, and Fernando Rios-Avila. (2013) "Identifying Factors behind the Decline in the Labor Force Participation Rate." *Business and Economic Research*, 3, 257–75.

Howard, Gregory, Robert Martin, and BethAnne Wilson. (2011) "Are Recoveries from Banking and Financial Crises Really So Different?" International Finance Discussion Paper 1037.

Juhn, Chinhui. (1992) "Decline of Male Labor Market Participation: The Role of Declining Market Opportunities." *Quarterly Journal of Economics*, 107, 79–121.

Kudlyak, Marianna, and Felipe Schwartzman. (2012) "Accounting for Unemployment in the Great Recession: Nonparticipation Matters." Federal Reserve Bank of Richmond Working Paper No. 12-04.

Leonard, Jonathan. (1979) "The Social Security Disability Program and Labor Force Participation." National Bureau for Economic Research Working Paper No. 392.

Macunovich, Diane. (2010) "Reversals in the Patterns of Women's Labor Supply in the United States, 1977–2009." *Monthly Labor Review*, 133, 16–36.

Mincer, Jacob. (1966) "Labor Force Participation and Unemployment: A Review of Recent Evidence." In *Prosperity and Unemployment*, edited by R. A. Gordon and M. S. Gordon, pp. 73–112. New York: Wiley.

Moffitt, Robert. (2012) "The Reversal of the Employment-Population Ratio in the 2000s: Facts and Explanations." Brookings Papers on Economic Activity, 43, 201–50.

Neiss, Katherine, and Edward Nelson. (2003) "The Real-Interest-Rate Gap as an Inflation Indicator." *Macroeconomic Dynamics*, 7, 239–62.

Orphanides, Athanasios. (2002) "Monetary Policy Rules and the Great Inflation." *American Economic Review*, 92, 115–20.

Pencavel, John. (1986) "Labor Supply of Men: A Survey. " In *Handbook of Labor Economics*, Vol. 1, edited by O. Ashenfelter and R. Layard, pp. 3–101. Amsterdam: Elsevier Science Publishers.

Sherk, James. (2012) "Not Looking for Work: Why Labor Force Participation Has Fallen During the Recession." *Backgrounder* Update Number 2722. Washington, DC: Heritage Foundation Center for Data Analysis.

Shimer, Robert. (2008) "The Probability of Finding a Job." *American Economic Review*, 98, 268–73.

Smith, Christopher. (2011) "Polarization, Immigration, Education: What's behind the Dramatic Decline in Youth Employment?" Finance and Economics Discussion Series 2011-41, Board of Governors of the Federal Reserve System.

Taylor, John, and John C. Williams. (2010) "Simple and Robust Rules for Monetary Policy." Federal Reserve Bank of San Francisco Working Paper 2010-10.

Toossi, Mitra. (2007) "Labor Force Projections to 2016: More Workers in Their Golden Years." *Monthly Labor Review*, 130, 33–52.

Van Zandweghe, Willem. (2012) "Interpreting the Recent Decline in Labor Force Participation." *Economic Review*, *Federal Reserve Bank of Kansas City*, 5–34.

Veracierto, Marcelo. (2008) "On the Cyclical Behavior of Employment, Unemployment and Labor Force Participation." *Journal of Monetary Economics*, 55, 1143–57.

Woodford, Michael. (2003) *Interest and Prices*. Princeton, NJ: Princeton University Press.