

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

Christopher J. Erceg and Andrew T. Levin

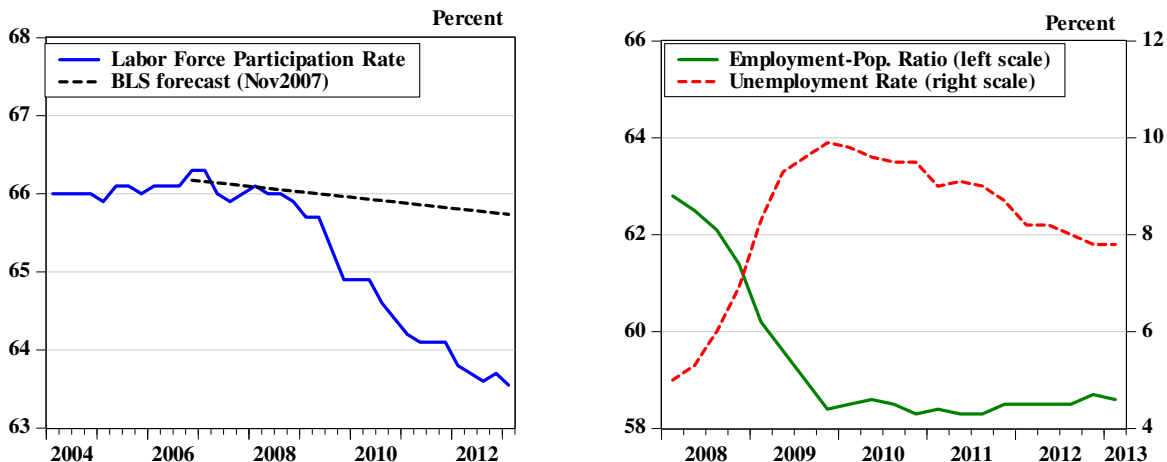
April 9, 2013

Abstract

In this paper, we provide compelling evidence that cyclical factors account for the bulk of the post-2007 decline in the U.S. labor force participation rate. We then proceed to formulate a stylized New Keynesian model in which labor force participation is essentially acyclical during "normal times" (that is, in response to small or transitory shocks) but drops markedly in the wake of a large and persistent aggregate demand shock. Finally, we show that these considerations can have potentially crucial implications for the design of monetary policy, especially under circumstances in which adjustments to the short-term interest rate are constrained by the zero lower bound.

*The authors are currently on leave from the Federal Reserve Board as visitors in the research department at the International Monetary Fund and may be contacted as follows—Erceg: phone 202-623-4564, email cerceg@imf.org; Levin: phone 202-623-9648, email alevin@imf.org. In the course of preparing this paper, we have appreciated very helpful comments and suggestions from George Akerlof, Olivier Blanchard, Mary Daly, Jeff Fuhrer, Bart Hobijn, and Prakash Loungani. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the International Monetary Fund, the Board of Governors of the Federal Reserve System, or any other person associated with the Federal Reserve System.

Figure 1: Labor Market Developments in the Wake of the Great Recession



Note: In the left panel, the solid line denotes the actual evolution of the labor force participation rate from 2004:Q1 to 2013:Q1, and the dashed line denotes the projected path for the decade ending in 2016 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 16 years and older (left scale), and the dashed line denotes the civilian unemployment rate (right scale).

1. Introduction

A longstanding and well-established fact in labor economics is that the labor supply of prime-age and older adults has been essentially acyclical throughout the postwar period, while that of teenagers has been moderately procyclical; cf. Mincer (1966), Pencavel (1986), and Heckman and Killingsworth (1986). Consequently, macroeconomists have largely focused on the unemployment rate 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; cf. 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 Figure 1, the labor force participation rate

¹For 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.

has fallen about 2-1/2 percentage points – a striking contrast to the modest decline that was projected by the Bureau of Labor Statistics in November 2007 just prior to the onset of the recession. Given the dropoff in labor force participation, the employment-to-population ratio has remained close to its pre-crisis trough even as the unemployment rate has returned roughly halfway back from its peak.

Our paper provides compelling empirical evidence that cyclical factors account for the bulk of the recent decline in the labor force participation rate (henceforth 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 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 2 percentage points in the LFPR for prime-age adults (that is, 25 to 54 years old). We define the participation gap as the deviation of the LFPR from its trend path as implied solely by demographic considerations, and we find that as of early 2013 this gap stood at around 2 percent—roughly the same magnitude as the unemployment gap (that is, the deviation of unemployment from its longer-run normal rate). Indeed, our analysis suggests that the participation gap and the unemployment gap each account for roughly half of the current employment gap, that is, the shortfall of the employment-to-population rate from its pre-crisis trend.

Our empirical analysis is broadly consistent with a number of other recent studies. Aaronson, Davis and Hu (2012) estimated statistical models for 44 demographic groups (based on age, gender, and educational attainment), incorporating birth cohort effects and other controls, and showed 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) found that cyclical factors accounted for 50 to 90 percent of the decline in the LFPR, depending on which measure of unemployment was used in constructing the filter. Sherk (2012) analyzed micro data from the Current Population Survey (CPS) and found that demographic factors only accounted for one-fifth of the post-recession decline in LFPR. Finally, Hotchkiss and Rios-Avila (2013) estimated a behavioral model of labor supply using CPS micro data and concluded that the decline in LFPR since the Great Recession was more than fully explained by the deterioration in labor market conditions.²

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 post-war 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.³ In normal recessions that are fairly transient in duration, the employment gap is largely driven by sharp but short-lived movements in the unemployment rate: labor force participation doesn't move much due to adjustment costs. However, a deep and protracted recession may eventually cause a sizeable decline in the LFPR. Importantly, to the extent that labor force participation responds very gradually to the unemployment rate, labor force participation may remain well below trend even as the economy begins recovering and the unemployment gap closes.

A second key feature of our model is 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 much less influence than the unemployment rate quantitatively. An important implication of this modified Phillips Curve is that inflation

²See Daly, Elias, Hobijn, and Jorda (2012) for further analysis and discussion.

³We adopt an alternative decentralization of the workhorse New Keynesian model so that changes in aggregate demand operate along the extensive margin of unemployment rather than the intensive margin of hours worked.

would remain below baseline following a recession even after the unemployment gap is closed, at least while the participation gap remains negative.⁴

The possibility that deep recessions may generate large cyclical swings in labor force participation has important implications for monetary policy design: should monetary policy respond to the cyclical component of labor force participation in some way, or focus more exclusively on the unemployment rate as suggested by a large literature focused on the Great Moderation? To address this question, we use our model to analyze several alternative monetary policy strategies against the backdrop of a deep recession that leaves labor force participation well below its long-run potential level. In our simulations, the deep recession reflects that the zero lower bound precludes monetary policy from lowering policy rates enough to offset a negative aggregate demand shock; once the shock dies away sufficiently, policy responds according to a non-inertial Taylor rule.

A key result of our analysis is that a monetary policy can induce a more rapid closure of the participation gap through allowing the unemployment rate to overshoot its long-run natural rate (i.e., unemployment falls below the natural rate). Quite intuitively, keeping unemployment persistently low draws cyclical non-participants back into labor force more quickly. Given that the cyclical non-participants exert some downward pressure on inflation, some overshooting 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 tradeoff 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 overshooting of the unemployment rate and inflation as the economy recovers. Even so, we stress that the

⁴Equivalently, 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.

overshooting of the unemployment rate in our analysis occurs under a *non-inertial* Taylor rule, and reflects that the participation gap remains sizeably negative even after 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), and consequentially, the monetary rules developed for the Great Moderation period may have to be adapted to account for broader measures of slack. Of course, as we emphasize below, our model does not incorporate a number of factors that may influence the tradeoff between employment gap and inflation stabilization, and hence the results of our analysis should certainly be taken with a grain of salt in contemplating the practical design of monetary policy strategy.

The remainder of this paper is organized as follows. Section 2 presents our empirical analysis. Section 3 describes our New Keynesian model with labor force participation, discusses the calibration, and characterizes model behavior under a benchmark policy rule. Section 4 considers the implications of alternative monetary policy strategies, while Section 5 concludes.


2. Empirical Analysis

In this section, we document the essentially acyclical behavior of the LFPR during “normal times” (that is, the postwar period from 1948 to 2007), and then we examine the extent to which the post-crisis decline in the LFPR is attributable to demographic vs. cyclical factors.

2.1. Labor Force Participation during “Normal Times”

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

Figure 2: Demographic Trends in Labor Force Participation, 1948 to 2007



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.

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 demographic groups over the period from 1948 to 2007.⁵ Of course, many volumes have been written about postwar trends in U.S. labor supply.⁶ Therefore, we will simply highlight a few broad features that will be salient for our subsequent analysis.

Prime-Age Males (25 to 54 years) comprised about 37 percent 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 Social Security Disability Insurance (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.⁷ However, those trends appear to have subsided over the half-decade prior to the Great Recession; that is,

⁵Appendix Figure A1 depicts the postwar evolution of the composition of the labor force by age and gender.

⁶See 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.

⁷See Leonard (1979), Juhn (1992), Gruber (2000), Bound and Waidmann (2002), Autor and Duggan (2003, 2006), and Moffitt (2012).

the LFPR for prime-age males was stable at around 90.5 percent from 2003 to 2007.

Prime-Age Females (25 to 54 years) comprised about 31 percent 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 percent—more than twice as high as in 1948. Interestingly, micro data indicates 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.⁸

Youths (aged 16 to 24 years) comprised about 15 percent 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 post-secondary education. The 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.⁹

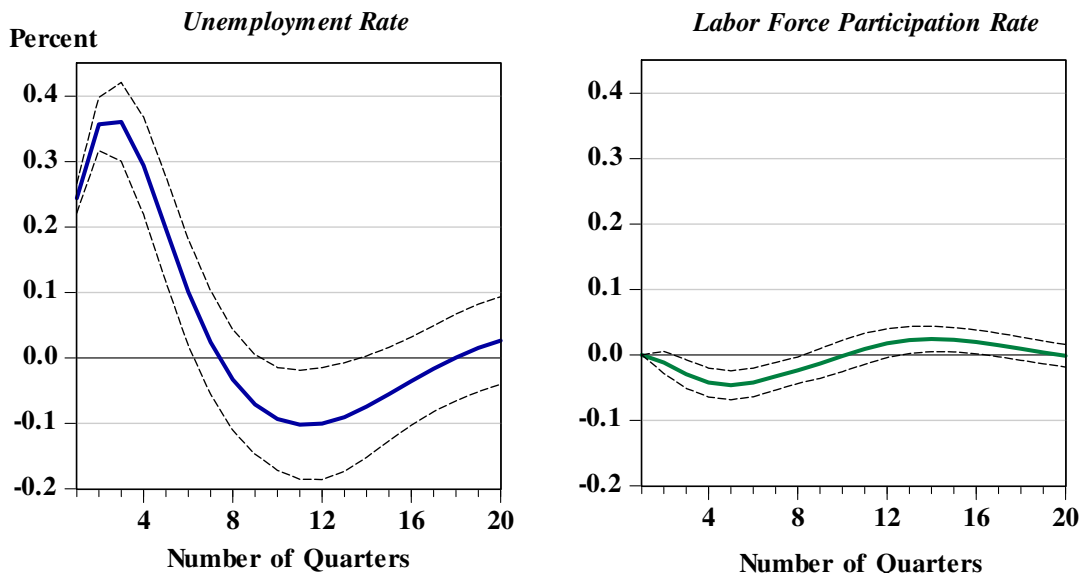
Older Adults (aged 55 years and above) comprised 17 percent of the labor force in 2007. The LFPR for older males declined gradually over the course of four decades to a trough of around 40 percent in the mid-1990s, while the LFPR for older females picked up slightly in the 1950s and then remained fairly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

Figure 2 also seems broadly consistent with the conventional wisdom that labor force participation has been practically acyclical throughout the postwar period. Indeed, even the steep recession of 1981-82 is essentially invisible in this figure. At any rate, any cyclical fluctuations in these time series have evidently been swamped by the demographic trends. Thus, we now focus more specifically on characterizing the cyclical behavior of labor force

⁸See Goldin (2006), Heim (2007), and references therein.

⁹See Smith (2011) and references therein.

Figure 3: The Cyclical Dynamics of the LFPR During "Normal Times"



Note: : This figure depicts the impulse responses from a bivariate VAR estimated using quarterly data on the HP-filtered unemployment rate and LFPR over the period 1948:Q1 to 2007:Q4. In each panel, the solid line denotes the response of the specified variable to an orthogonalized unemployment rate shock of one standard deviation, and the dashed lines depict 95 percent confidence bands based on 10,000 Monte Carlo replications.

participation, applying the HP filter (with a smoothing parameter of 1600) to quarterly data on the LFPR and the unemployment rate over the period 1948:Q1 to 2007:Q4. The standard deviation of the HP-filtered LFPR series is 0.24 percent—much smaller than the standard deviation of 0.8 percent for the HP-filtered unemployment rate. Nonetheless, the correlation coefficient of -0.35 is statistically significant and indicates that labor force participation does indeed tend to move systematically, though modestly, during normal business cycles.

To shed more light on these cyclical dynamics, we estimate a bivariate vector autoregression (VAR) using the filtered unemployment and LFPR series, and then we compute the impulse response of each variable to an orthogonalized unemployment rate shock of one standard deviation.¹⁰ As shown in Figure 3, the peak decline in the LFPR is only -0.04

¹⁰The impulse responses shown here are computed from a bivariate VAR with lags of order 1, 2 and 7 (chosen using an exclusion test procedure), but the results are not sensitive to alternative choices for the lag order. We have also confirmed that the results are robust to estimating the VAR using the raw data rather than the HP-filtered series, although in that case the orthogonalized unemployment shock also embeds shifts in the natural unemployment rate.

percent, an order of magnitude smaller than that of the peak for the unemployment rate.¹¹ In effect, the cyclical movement in the LFPR would remain practically invisible—less than a tenth of a percentage point—even if the shock were twice as large.

Figure 3 also highlights the extent to which the LFPR exhibits even more inertia than the unemployment rate. In particular, the peak response for the unemployment rate occurs within a couple of quarters, whereas the peak response for LFPR takes roughly twice as long. Indeed, at a horizon of 10 quarters, the unemployment rate has declined all the way to its cyclical trough, whereas the LFPR is still just edging up near zero.

Evidently, the highly inertial behavior of the LFPR accounts for its practically acyclical behavior during normal times; i.e., even a sharp recession (like that of 1981-82) hasn't had much impact on the LFPR because the subsequent recovery has been V-shaped. Nonetheless, these same inertial dynamics also imply that the LFPR might well exhibit a much larger cyclical response under circumstances in which a deep recession was followed by a very slow recovery and hence a protracted period of relatively high unemployment.

2.2. Labor Force Participation Since the Great Recession

As we have already noted above in discussing Figure 1, the LFPR declined by about $2\frac{1}{2}$ percentage points over the five-year period ending in 2013:Q1. Thus, the crucial task is to gauge the extent to which that decline reflects cyclical vs. demographic factors. To address this issue, we first consider the post-crisis evolution of the LFPR for specific demographic groups, and then we perform regression analysis using data on LFPR and unemployment for U.S. states. Finally, we explore the issue of hysteresis by examining the extent to which macroeconomic conditions may have contributed to a greater incidence of workers becoming SSDI beneficiaries and effectively departing permanently from the labor force.

In our view, the labor force projections published by the BLS in November 2007 serve as

¹¹ As shown in Appendix Figure A2, the impulse response of the LFPR for prime-age adults is quantitatively similar to that shown in Figure 3 for the aggregate LFPR, whereas the LFPR for youth exhibits much greater cyclicity and markedly less persistence, and the LFPR for older adults is truly acyclical.

an invaluable resource in assessing the influence of demographic factors on the subsequent decline in the LFPR. In making such projections, BLS staff consider detailed demographic groups using state-of-the-art statistical procedures in conjunction with micro data from the Current Population Survey (CPS) and various 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 one month later, and hence those BLS projections effectively encompassed all of the pre-recession 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 labor force participation rate 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

Table 1: Demographic Factors and the Recent Evolution of the LFPR

Demographic Group	Population Share			Labor Force Participation Rate		
	2007	Change (2007 to 2012)		2007	Change (2007 to 2012)	
	<i>Actual</i>	<i>Projection</i>	<i>Actual</i>	<i>Actual</i>	<i>Projection</i>	<i>Actual</i>
<i>16 to 24 yrs</i>	16.1	-0.9	-0.2	59.4	-0.9	-4.5
<i>25 to 54 yrs</i>	54.2	-2.0	-3.1	83.0	0.3	-1.5
<i>55 to 64 yrs</i>	14.0	1.3	1.8	64.0	1.5	0.5
<i>65 and older</i>	15.6	1.7	1.6	15.4	2.6	3.1
<i>Total</i>	100	0	0	66.1	-0.3	-2.4

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

adults would continue rising notably over coming years, consistent with the trends that had prevailed since the mid-1990s (as noted above 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, 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. 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; i.e., 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 above. 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 still only a fourth as large as its actual decline of $2\frac{1}{2}$ percent.

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.


2.3. Distinguishing the Role of Cyclical vs. Structural Factors

The preceding analysis underscores a crucial question: Why did the LFPR for prime-age adults decline by nearly two percentage points from 2008:Q1 to 2013:Q1, given that the rate for this demographic group had been essentially stable over the 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*—though it should be noted that such a rapid pace of decline would have been unprecedented for both prime-age males and females, at least since the end of WWII. The obvious 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*. Such an interpretation would seem to conflict with the essentially acyclical characteristic of prime-age LFPR throughout the postwar period, but of course the depth of the Great Recession and the sluggishness of the subsequent recovery also stand in marked contrast to all other postwar business cycles.

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.¹² 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,

¹²Clark and Summers (1982) explored the relationship between the *levels* of LFPR and unemployment using cross-sectional data from the 1970 U.S. Census. Aaronson, Fallick, Figura, Pingle, and Wascher (2006) analyzed a panel of annual data for U.S. states over the period 1990 to 2005 and found a statistically significant relationship between the LFPR and the contemporaneous unemployment rate, using lagged unemployment rates as instruments.

Figure 4: State-Level Data on Unemployment and LFPR for Prime-Age Adults



Note: The vertical axis refers to the change in the labor force participation rate for prime-age adults (25 to 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.

many other states experienced practically catastrophic outcomes: During 2008-2009, 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 vs. structural factors, we estimate the

following linear regression using ordinary least squares:

$$\Delta LFPR_i = \alpha + \beta \Delta UNEMP_i + \varepsilon_i \quad (1)$$

where $\Delta LFPR_i$ denotes the change in the LFPR for prime-age adults in state i over the period 2007 to 2012, 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 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 provides no support whatsoever for structural interpretations of the drop in prime-age LFPR. In effect, the state-level data indicates that the aggregate decline in prime-age LFPR since 2007 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.¹³

¹³In 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-specific and state-specific effects.

Table 2: Regression Analysis of State-Level Data for Prime-Age Adults

Explanatory Variable	All U.S. States <i>(including Washington DC)</i>			41 Larger States <i>(excluding Nevada)</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Intercept</i>	-0.40 <i>(0.42)</i>	-0.41 <i>(0.42)</i>	—	-0.18 <i>(0.55)</i>	-0.18 <i>(0.56)</i>	—
$\Delta Unemp$ <i>(2007-2010)</i>	-0.30 <i>(0.08)</i>	-0.27 <i>(0.11)</i>	-0.39 <i>(0.03)</i>	-0.33 <i>(0.10)</i>	-0.31 <i>(0.14)</i>	-0.37 <i>(0.03)</i>
$\Delta Unemp$ <i>(2010-2012)</i>	—	0.10 <i>(0.24)</i>	—	—	0.08 <i>(0.26)</i>	—

Note: 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. Equations 2 and 5 also include the change in the prime-age unemployment rate from 2010 to 2012. Equations 1, 2, 4 and 5 also include a regression intercept. Standard errors are adjusted for heteroskedasticity (White 1980) and are given in parentheses.

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 DSGE model described in Section 3 below.

We have also confirmed that the regression results are robust to excluding potentially influential observations from the sample. In particular, as evident from Figure 4, 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 panel 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. Notably, these results are virtually identical to those reported in the left panel of the table.

2.4. Gauging the Potential for Hysteresis

Having concluded that the post-2007 decline in LFPR is mainly attributable to the Great Recession and its aftermath, we now turn to the issue of hysteresis. More specifically, to the extent that persistent shortfalls in labor demand have caused a large number of prime-age adults to leave the labor force, what is the likelihood that those individuals will remain permanently out of the labor force vs. re-engaging in the labor market once jobs become more readily available? Similarly, given the sharp increase in the number of youths who have refrained from entering the labor force since the onset of the Great Recession, what are the prospects that they will join the labor force over coming years? Of course, no definitive answers to these questions can be provided at this stage, but it does seem worthwhile to examine the available evidence on several specific factors, namely, educational enrollment, permanent disability, and retirement. Thus, we now consider the relevance of these factors for each major age group:

Youths (16 to 24 years). Monthly data from the CPS indicates that the educational enrollment rate for this demographic group jumped four percentage points from 2007 to 2011, essentially accounting for nearly all of the decline in the youth LFPR. By contrast, there has been practically no change in the incidence of SSDI beneficiaries under age 25.

(Only a tiny fraction of youths not in the labor force describe themselves as "retired.")¹⁴

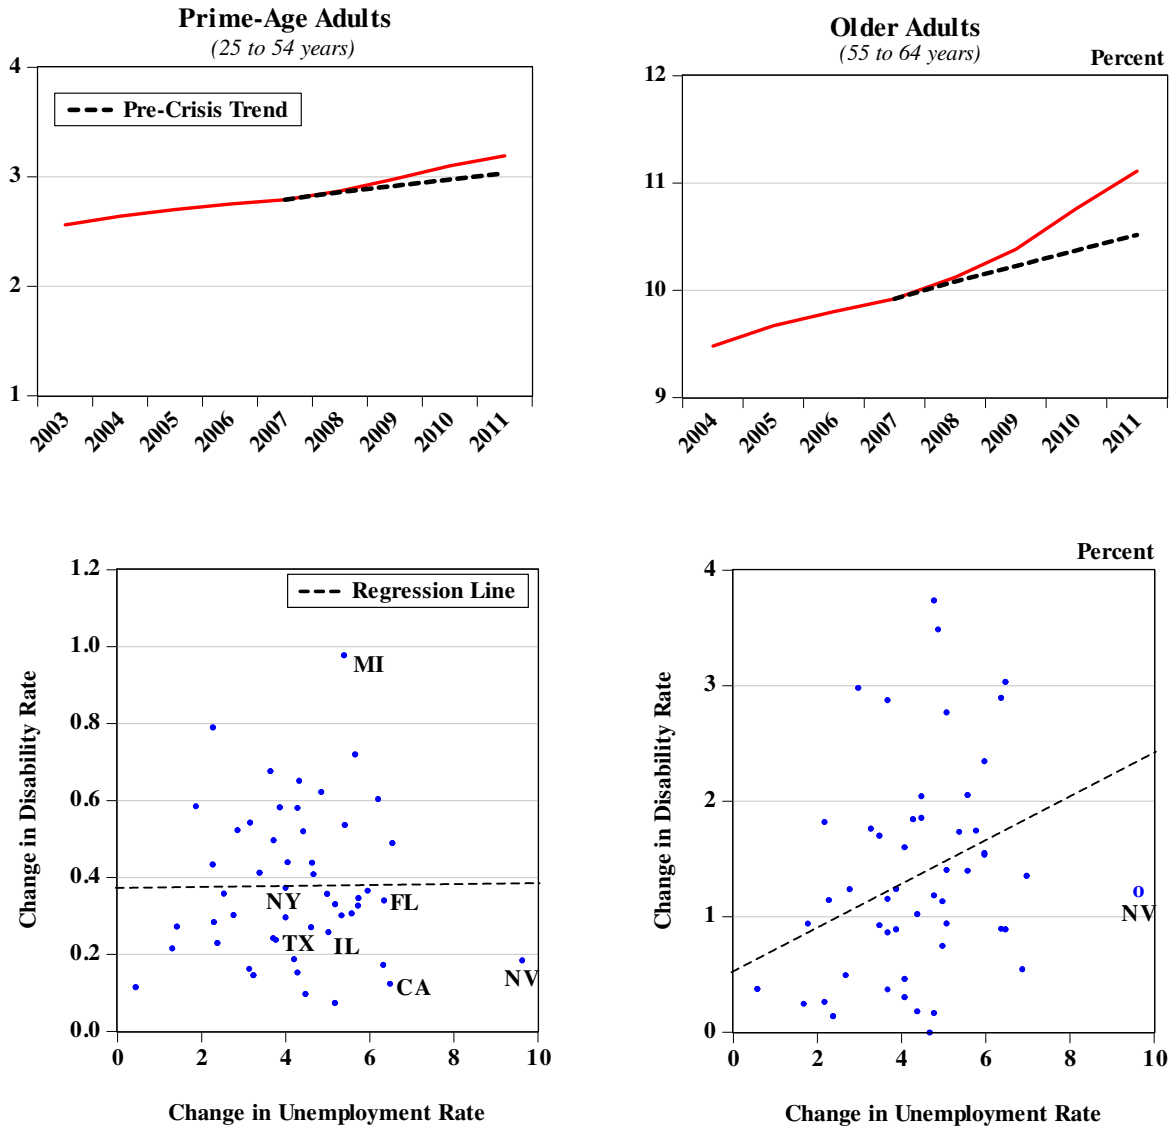
Prime-Age Adults (25 to 54 years). CPS data indicates that the educational enrollment rate of prime-age adults rose about 0.6 percentage points from 2007 to 2011, accounting for nearly a third of the overall decline in the prime-age LFPR. Over that period, there was also a slight increase—about a tenth of a percentage point—in the number of prime-age adults describing themselves as "retired." As shown in the upper-left panel of Figure 5, 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 pre-crisis 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 primarily reflect changes in the screening process.¹⁵ Nonetheless, these factors only explain about half of the post-2007 decline of about 2 percentage points in the prime-age LFPR. Evidently, the remainder represents roughly a million individuals who have given up searching for a job and instead are engaged in other activities such as child care, home renovation projects, etc.

Older Adults (55 years and above). As shown above in Table 1, the LFPR for adults 55 to 64 years old has risen only modestly since 2007—a full percentage point lower in 2012 than the BLS had projected as of November 2007. The upper-right panel of Figure 5 points to one potentially important factor, namely, an acceleration in the number of workers aged 55 to 64 years old who have become SSDI beneficiaries. Moreover, as shown in the

¹⁴According to the SSA's latest annual report, there was a total of 167,651 SSDI beneficiaries aged 16 to 24 years old as of December 2011, an increase of about 34,000 since December 2007—a miniscule change relative to the total population of about 38 million for 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 permanently disabled prior to age 22 and whose eligibility was based on the work experience of the person's parents).

¹⁵As shown in the figure, the regression line fitted to this data is nearly horizontal, reflecting the fact that the slope coefficient is close to zero (and not statistically significant).

Figure 5: The Incidence of Permanent Disabilities



Note: In the upper-left panel, the solid line depicts the number of prime-age workers receiving Social Security disability benefits (that is, beneficiaries aged 25 to 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 to 54 years, and the dashed line denotes the projected values from fitting a linear trend over the period 2003 to 2007. The upper-right panel provides corresponding information for older adults, that is, the number of disabled workers 55 to 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.

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, 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 pre-crisis 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 pre-crisis trend, at least partly because the increased incidence of SSDI beneficiaries aged 55 to 64 has been offset by the greater incidence of working adults aged 65 and above.

In summary, this analysis suggests that the degree of hysteresis in the aggregate LFPR is likely to be moderate over coming years. A substantial proportion of the recent drop in the LFPR reflects individuals who have chosen to expand their educational attainment in ways that would presumably enhance their employment opportunities over time. Moreover, some prime-age adults would presumably consider rejoining the labor force if they perceived a substantial improvement in the prospects for finding a job. Unfortunately, in light of previous empirical work, it seems likely that those older adults who have recently become SSDI beneficiaries will remain permanently out of the labor force, although even that outlook might be affected by how labor market conditions evolve over coming years.

2.5. Assessing 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 deviation between civilian employment (E_t) and its maximum sustainable level (E_t^*) as a share of the civilian noninstitutionalized population (POP_t). In particular, the employment gap satisfies the

following relationship:

$$EGAP_t = (1 - u_t^*) (LFPR_t^* - LFPR_t) + (u_t - u_t^*) LFPR_t^* + (u_t - u_t^*) (LFPR_t^* - LFPR_t) \quad (2)$$


where $LFPR_t$ refers to the actual labor force participation rate, $LFPR_t^*$ denotes its level based solely on demographic trends, u_t refers to the unemployment rate, and u_t^* denotes its longer-run sustainable rate. To facilitate the discussion, it is convenient to define the participation gap $LFGAP_t = LFPR_t^* - LFPR_t$, 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 hence will generally be negligible compared with 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:

$$EGAP_t \simeq (1 - u_t^*) LFGAP_t + LFPR_t^* (u_t - u_t^*) \quad (3)$$

where 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 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 Congressional Budget Office's 10-year-ahead unemployment rate forecast as the value of u_t^* .

As shown in Figure 6, the employment gap rose sharply during the Great Recession to a peak above 4 percent 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 pre-crisis trend. This figure also underscores the excruciatingly sluggish nature of the recovery. The employment gap has

Figure 6: The Magnitude and Composition of the Employment Gap



Note: The solid line denotes the employment gap, that is, the percent deviation between U.S. civilian employment and its maximum sustainable level. As discussed in the text, this gap can be expressed as the sum of two components: (1) the participation gap weighted by the proportionality factor $(1-u^*)$; and (2) the unemployment gap weighted by the proportionality factor $LFPR^*$. These components are denoted by the short-dashed and long-dashed lines, respectively.

only narrowed modestly over the past several years, because the steady widening of the participation gap has been comparable in magnitude to the decline in the unemployment gap.¹⁶ Indeed, our analysis suggests that the participation gap and the unemployment gap each account for roughly half of the current level of the employment gap.

¹⁶The 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 Appendix Figure A5, such individuals only comprise a modest fraction of the overall participation gap.

3. A New Keynesian Model with Labor Force Participation

In this section we describe a New Keynesian model with endogenous labor force participation. In the spirit of the literature on home production, labor supply decisions are made by a representative household that must choose how to allocate labor between market and home (or non-market) production. However, to account for unemployment in the market sector, it is convenient to assume that the representative household allocates “families” between the two sectors. 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 labor force participation rate. As in the home production literature, the labor force participation rate depends on the relative return to working in the market relative to the home production sector; however, we also introduce adjustment costs to capture the empirically realistic feature that flows between the two sectors appear to respond very gradually.

If labor force participation were held fixed, our model would be observationally equivalent to the workhorse three equation New Keynesian model. The only difference would be that the workhorse New Keynesian model specifies labor market variables along the intensive margin of hours worked, whereas we adopt a decentralization (described below) along the extensive margin of the unemployment rate (so that the “hours gap” in the former is reinterpreted as reflecting a gap between the employment rate H_{mt} of the fixed labor force and its level under fully flexible prices H_{mt}^{pot}).¹⁷ With endogenous labor force participation, marginal cost in the Phillips Curve relation turns out to depend not only on the unemployment gap, but also on the cyclical labor force participation gap $L_{Ft} - L_{Ft}^{pot}$.

¹⁷This formulation attempts to capture in a simple way how weak aggregate demand leads to unemployment, which is the quintessential feature of Keynesian models.

3.1. Households

To rationalize unemployment in the market sector, we assume that each family is itself comprised of a continuum of members distributed on the unit interval. Similar to Gali (2011), these members can be regarded as ordered sequentially based on their disutility from working during the period, with individual h experiencing a 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 economywide 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:

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \epsilon_{dt+j} \left\{ \begin{array}{l} \frac{1}{1-\sigma_C} (C_{t+j} - \psi_c C_{t+j-1}^a)^{1-\sigma_C} - L_{Ft} \frac{\chi_0 H_{mt+j}^{1+\chi}}{1+\chi} + \frac{\gamma_G}{1-\sigma_G} L_{Ht}^{1-\sigma_G} \\ -0.5\phi_G \left(\frac{L_{Ft+j}}{L_{Ft+j-1}} - 1 \right)^2 L_{Ft+j} \end{array} \right\} \quad (4)$$

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. 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 $L_{Ft} \frac{\chi_0 H_{mt}^{1+\chi}}{1+\chi}$. 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.

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:

$$P_t C_t + B_{G,t} - (1 + i_{t-1}) B_{G,t-1} = W_t N_t + \Gamma_t \quad (5)$$

Here $B_{G,t}$ are purchases of discount 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 labor force participation rate L_{Ft} , i.e., $N_t = H_{mt} L_{Ft}$.

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

$$\frac{dV(H_{mt})}{dH_{mt}} = \chi_0 H_{mt}^\chi = \frac{W_t}{P_t} \lambda_t \quad (6)$$

For each working family, equation (6) 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 ag-

gregate demand is interpreted as reducing the employment rate relative to its flexible price level (or equivalently, boosting the unemployment rate above its potential rate), 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:

$$\left\{ \lambda_t \frac{W_t}{P_t} H_{mt} - V(H_{mt}) \right\} = \frac{dM(L_{ht})}{dL_{ht}} = (1 - L_{ft})^{-\sigma_G} \quad (7)$$

Equation (7) 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 (6), $V(H_{mt})$ may be expressed as $V(H_{mt}) = \lambda_t \frac{W_t}{P_t} \frac{H_{mt}}{1+\chi}$. Substituting into equation (7) and using (6) yields:

$$\left\{ \frac{\chi}{1+\chi} \lambda_t \frac{W_t}{P_t} H_{mt} \right\} = \frac{\chi}{1+\chi} \chi_0 H_{mt}^{1+\chi} = (1 - L_{ft})^{-\sigma_G} \quad (8)$$

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 below, 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:

$$\lambda_t = E_t \left(\beta \frac{\epsilon_{dt+1}}{\epsilon_{dt}} \right) \lambda_{t+1} \frac{1 + i_t}{1 + \pi_{t+1}} \quad (9)$$

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 $(\frac{P_t}{P_{t-1}} - 1)$.

3.2. 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:

$$Y_t(f) = K_t(f)^\alpha N_t(f)^{1-\alpha} \quad (10)$$

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

$$\frac{MC_t}{P_t} = \frac{W_t/P_t}{MPL_t} = \frac{W_t/P_t}{(1-\alpha)K^\alpha N_t^{-\alpha}}. \quad (11)$$

Each monopolistically competitive firm faces a downward-sloping demand curve of the form $Y_t(f) = \left[\frac{P_t(f)}{P_t}\right]^{\frac{-(1+\theta_p)}{\theta_p}} Y_t$, 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 ($P_t^{opt}(f)$) solves the usual problem:

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

where $\psi_{t,t+j}$ is the stochastic discount factor (the conditional value of future profits in utility units, i.e. $\beta^j \mathbf{E}_t \frac{\lambda_{t+j} P_{t+j}}{\lambda_t P_t}$). 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 re-optimize mechanically adjust their price in line with past inflation.

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

3.3. Log-Linearized Model

The key equations of the model are:

$$n_t - n_t^{pot} = (n_{t+1|t} - n_{t+1|t}^{pot}) - \frac{1 - \alpha}{\hat{\sigma}_c} (i_t - \pi_{t+1|t} - r_t^{pot}), \quad (12)$$

$$\pi_t - \nu\pi_{t-1} = \beta(\pi_{t+1|t} - \nu\pi_t) + \kappa_p mc_t, \quad (13)$$

$$n_t - n_t^{pot} = (l_{ft} - l_{ft}^{pot}) + (e_{rt} - e_{rt}^{pot}) = (l_{ft} - l_{ft}^{pot}) - (u_{rt} - u_{rt}^{pot}), \quad (14)$$

$$mc_t = \frac{w_t}{p_t} - MPL_t = \psi_e (u_{rt}^{pot} - u_{rt}) + \psi_l (l_{ft} - l_{ft}^{pot}), \quad (15)$$

$$l_{ft} - l_{ft}^{pot} = \mu_1 (l_{ft-1} - l_{ft-1}^{pot}) + \mu_2 (l_{ft+1|t} - l_{ft+1|t}^{pot}) - \phi_m (u_{rt} - u_{rt}^{pot}), \quad (16)$$

$$i_t = \max \left(-i, \gamma_\pi \pi_t + \gamma_n (n_t - n_t^{pot}) - \gamma_u (u_{rt} - u_{rt}^{pot}) + \varepsilon_{it} \right), \quad (17)$$

$$n_t^{pot} = \left(\frac{1 + \phi_m}{\chi + (1 + \phi_m)(\alpha + (1 - \alpha)\hat{\sigma}_c)} \right) [(1 - \hat{\sigma}_c)z_t + \hat{\sigma}_c \varepsilon_c \varepsilon_{ct}], \quad (18)$$

$$r_t^{pot} = \hat{\sigma}_c \left(1 - \frac{1}{\phi_{mc} \hat{\sigma}_c} \right) \left[(1 - \alpha)(n_{t+1|t}^{pot} - n_t^{pot}) + \varepsilon_c (\varepsilon_{ct} - \varepsilon_{ct+1|t}) \right] + \varepsilon_d (\varepsilon_{dt} - \varepsilon_{dt+1|t}), \quad (19)$$

where $\hat{\sigma}_c$, κ_p , and ϕ_{mc} are composite parameters defined as:

$$\hat{\sigma}_c = \frac{\sigma_c}{1 - \varepsilon_c}, \quad (20)$$

$$\kappa_p = \frac{(1 - \xi_p)(1 - \beta\xi_p)}{\xi_p}, \quad (21)$$

$$\phi_m = \frac{1 + \chi}{\sigma_G \left(\frac{l_f}{1 - l_f} \right)} \quad (22)$$

$$\psi_l = \alpha + (1 - \alpha)\hat{\sigma}_c, \quad \psi_e = \chi + \psi_l \quad (23)$$

All variables are measured as percent or percentage point deviations from their steady state level.¹⁸ The superscript ‘pot’ denotes the level of a variable that would prevail under completely flexible prices, e.g., y_t^{pot} is *long-run* potential output. For expositional simplicity,

¹⁸We use the notation $y_{t+j|t}$ to denote the conditional expectation of a variable y at period $t + j$ based on information available at t , i.e., $y_{t+j|t} = E_t y_{t+j}$.

the equations above abstract from habit persistence in consumption.

Equation (12) expresses the “New Keynesian” IS curve. The overall employment gap $n_t - n_t^{pot}$ depends inversely on the deviation of the real interest rate ($i_t - \pi_{t+1|t}$) from its potential rate r_t^{pot} , as well as on the expected employment gap in the following period. The interest sensitivity of the overall employment gap depends both on the household’s intertemporal elasticity of substitution in consumption $\frac{1}{\sigma_c}$, and directly on the labor share, with the latter reflecting that the overall employment gap is simply proportional to the output gap $y_t - y_t^{pot}$ (since $n_t = (1 - \alpha)y_t$). The overall employment gap should be regarded as capturing the “cyclical” component of the employment to population ratio: as seen from equation (14), it is comprised of both the the labor force participation rate gap ($l_{ft} - l_{ft}^{pot}$) and the employment **rate** gap ($e_{rt} - e_{rt}^{pot}$); or alternatively, of the labor force participation gap minus the unemployment rate gap ($u_{rt} - u_{rt}^{pot}$).

The price-setting equation (13) specifies the inflation rate π_t to depend both on expected inflation and marginal cost. Given the Calvo-Yun contract structure, the composite parameter κ_p depends inversely on the mean contract duration ($\frac{1}{1-\xi_p}$), while the quasi-difference specification for inflation reflects that we allow for some degree of dynamic indexation of price contracts.

A novel feature of the aggregate supply block is that marginal cost depends both on the unemployment rate gap and the labor force participation gap, as seen in equation (15). Our model implies that inflation falls relative to target if there is slack in either labor market resource gap, i.e., if either $u_{rt} > u_{rt}^{pot}$, or if $l_{ft} < l_{ft}^{pot}$. The different response coefficients ψ_e and ψ_l reflect that the participation rate 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 ($\alpha = 0$) in which marginal cost varies only in response to the marginal rate of substitution (MRS) of

working households, the log-linearized version of equation (6) is given by:

$$mc_t = MRS_t = \chi e_{rt} + \hat{\sigma}_c(c_t - \hat{\sigma}_C \varepsilon_c \varepsilon_{ct}) = -(\chi + \hat{\sigma}_c)(u_{rt} - u_{rt}^{pot}) + \hat{\sigma}_c(l_{ft} - l_{ft}^{pot}), \quad (24)$$

so that the participation gap enters only through the wealth effect term $\hat{\sigma}_c(c_t - \hat{\sigma}_C \varepsilon_c \varepsilon_{ct})$, and captures that lower participation reduces consumption and hence shifts out the wage schedule. Our model constrains the coefficient on the unemployment rate gap to exceed that on the participation gap, with the coefficients converging only in the limiting case in which the Frisch elasticity of labor supply becomes infinitely large ($\chi = 0$). Stepping beyond our model, it seems quite reasonable that a large negative cyclical participation gap would exert at least *some* downward pressures on wages and marginal cost, especially to the extent that it reflected a falloff in the participation rate of relatively younger and geographically mobile members of the population.

Equation (16) describes how the labor force participation gap responds dynamically to the unemployment rate gap. The long-run response hinges on the parameter ϕ_m , which itself depends on the relative degree of curvature of the disutility of work schedule to 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. The general form of equation (16) implies that both current and future unemployment rate gaps may affect the response of current labor force participation. However, we mainly focus on the specification with external adjustment costs which implies that labor force participation follows a simple AR(1) (with the lag coefficient μ_1 , and the forward lead $\mu_2 = 0$). This specification has the appeal of simplicity, and seems reasonably plausible empirically.

In our baseline model, we assume that the policy rate i_t follows a simple Taylor rule subject to the zero lower bound (equation 17). In addition to inflation, the policy rate is assumed to react either to the unemployment rate gap $u_{rt} - u_{rt}^{pot}$, to the overall employment

gap $n_t - n_t^{pot}$, or possibly both. The policy rule is assumed to be non-inertial, 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 lagged state variable. We also consider optimal policies with full commitment in Section 4.

3.4. 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 $\pi = .005$; this implies a steady state interest rate of $i = .01$ (i.e., four percent 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 $\frac{1}{\chi} = 0.2$. These parameter values imply that $\psi_l = 1$ and $\psi_e = 6$ in the marginal cost expression (15), so that a one percentage point rise in the participation gap boosts marginal cost by 1/6th as much as a one percentage point fall in the unemployment gap. The price contract duration parameter ξ_p is set to 0.985, implying a very low degree of responsiveness of inflation to marginal cost of .0003. Given the absence of wage rigidities, a low sensitivity is important for generating plausible variation in inflation in response to deep recessions similar in magnitude to the Great Recession; but we also consider shorter contract durations in sensitivity analysis. Conditional on these parameters, 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 (16) equal 0.97 and 0.06, respectively; thus, labor force participation responds very slowly to the employment rate. Finally, the calibration of the monetary rule is described below.

3.5. 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 (17) 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 100 basis points at impact. The shock depresses overall employment by about 0.7 percent at impact, with habit persistence in consumption accounting for the slightly hump-shaped response.¹⁹ Importantly, the substantial near-term fall in the employment rate is driven almost entirely by a rise in the unemployment rate (not shown, but is given by the gap between the employment rate and labor force participation). 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:

$$l_{ft} = \mu_1 l_{ft-1} - \phi_m u_{rt}, \quad (25)$$

labor force participation can exhibit a large eventual drag on overall 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 three years (with the shocks having a persistence of 0.9). About half of the employment rate 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 recession. Moreover, as the economy recovers and the unemployment rate starts to fall (note that the gap between the lines begins narrowing after twelve quarters), labor force

¹⁹The implied peak response of output of about 0.5 percent is well within the range typically estimated in VAR studies.

participation continues to decline and accounts for progressively more of the employment gap. We consider below 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 trendlike deterioration in employment.

4. The Design of Monetary Policy Rules

We begin our analysis by considering tradeoffs in a simple static framework, and then proceed to model simulations.

4.1. Policy Tradeoffs

Our model with endogenous labor force participation is formally identical to the simple New Keynesian model in two special cases. One case occurs if labor force participation responds immediately 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 overall employment gap $n_t - n_t^{pot}$ in this case, marginal cost and hence inflation also depend only on $n_t - n_t^{pot}$. Thus, monetary policy would face no tradeoff between stabilizing inflation and the overall employment gap; and it would be immaterial whether monetary policy responded to the unemployment gap or overall employment gap. Similarly, there is no stabilization tradeoff if $\psi_l = \psi_e$. In this case, the labor force participation gap has an equal-sized effect on marginal cost as the unemployment gap, and hence marginal cost again depends only on the overall employment gap. In our model, this occurs as the Frisch elasticity of labor supply tends toward infinity.

Conversely, our model has implications that can differ substantively from that of the

New Keynesian model in the more plausible case in which labor force participation responds gradually to the unemployment rate – consistent with our previous empirical analysis – and $\psi_e > \psi_l$ (so that the unemployment rate gap matters relatively more for inflation determination). Several implications are particularly important. First, in the wake of a negative shock that depresses labor force participation – say because the zero bound keeps monetary policy from cutting interest rates enough – a policy oriented toward stabilizing the unemployment rate u_{rt} at its *long-run* potential of u_{rt}^{pot} is consistent with a fall in inflation below target. This can be inferred from equation (15), and reflects that a negative labor force participation gap puts some downward pressure on marginal cost even if $u_{rt} = u_{rt}^{pot}$. A second related implication is that keeping inflation stable in the wake of a persistently negative labor force participation gap requires allowing the unemployment rate to fall below its long-run potential rate for some time; as we discuss below, the short-run natural rate of unemployment may be regarded as falling in response to the “deflationary” pressure associated with the cyclical fall in labor force participation. Finally, a policy of keeping overall employment n_t close to its long-term natural rate of n_t^{pot} puts upward pressure on marginal cost and inflation. Hence, monetary policy faces a tradeoff between stabilizing the overall employment gap and inflation.

Figure 8 is helpful for providing intuition for these results by showing how the unemployment rate and marginal cost (hence inflation) are determined in a simple static setting. The long-run equilibrium is at point A, where the downward-sloping real wage schedule intersects the MPL (we again assume $\alpha = 0$ for simplicity). Wages and hence marginal costs rise as one moves northwest along this schedule since a lower unemployment rate boosts the real wage demanded by workers. The dashed line shows that the wage schedule would shift downward in response to a recessionary shock that lowered last period’s participation gap. Given that the marginal cost schedule may be expressed (using equations (24) and (16)) as:

$$mc_t = MRS_t = -(\chi + \hat{\sigma}_c(1 + \mu))(u_{rt} - u_{rt}^{pot}) + \hat{\sigma}_c\mu_1(l_{ft-1} - l_{ft}^{pot}), \quad (26)$$

the lower initial labor force participation shifts out labor supply for any given unemployment rate – an adverse wealth effect – which puts downward pressure on marginal costs. To stabilize marginal cost and hence inflation, monetary policy has to allow the unemployment rate u_{rt} to temporarily fall to point C, which is below its long-run level of u_{rt}^{pot} . Equivalently, the “short-run” natural rate of unemployment – defined as the level which keeps marginal cost at zero (relative to steady state) in equation (26) – falls by:

$$u_{rt}^{SR,pot} = u_{rt}^{pot} + \frac{\hat{\sigma}_c \mu_1}{\chi + \hat{\sigma}_c(1 + \mu)_1} (l_{ft-1} - l_{ft}^{pot}), \quad (27)$$

Conversely, if monetary policy targeted the unemployment rate at its *long-run* natural rate of u_{rt}^{pot} , marginal cost and inflation would fall below target (until the participation gap gradually reverted back to its long-run level). Intuitively, because those “outside the labor force” put downward pressure on the wages of workers, stabilizing the unemployment rate at its long-run level (once it becomes feasible to do so) is consistent with below-target inflation, and a slow recovery in labor force participation.

We next consider the consequences of a policy of setting overall employment n_t equal to n_t^{pot} (again, once it become feasible). As it turns out, this policy would imply upward pressure on marginal cost and inflation as illustrated by point D in Figure 8. To see this, note that the employment gap may be expressed:

$$n_t - n_t^{pot} = -(1 + \mu)(u_{rt} - u_{rt}^{pot}) + \mu_1(l_{ft-1} - l_{ft}^{pot}), \quad (28)$$

A policy that fully stabilized the unemployment rate at its short-term potential rate of $u_{rt}^{SR,pot}$ – and hence kept inflation at target – would imply $n_t < n_t^{pot}$, as can be seen by setting $u_{rt} = u_{rt}^{SR,pot}$ into equation (28) and substituting in equation (27):

$$n_t - n_t^{pot} = \mu_1 \left(1 - \frac{\hat{\sigma}_c(1 + \mu)}{\chi + \hat{\sigma}_c(1 + \mu)_1} \right) (l_{ft-1} - l_{ft}^{pot}) \quad (29)$$

Thus, assuming labor force participation is below its long-run level, inflation stabilization

requires keeping total employment below its long-run level. Moreover, a policy aimed at targeting the total long-run employment gap n_t^{pot} would boost marginal cost and hence inflation.

4.2. Alternative Policy Rules

Our static analysis 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 tradeoff in pursuing a more aggressive policy of stabilizing the overall employment gap. We next illustrate these implications using dynamic simulations of our model under both a non-inertial Taylor rule and under an “optimal” policy rule derived from minimizing an intertemporal loss function. The section concludes with robustness exercises that show how the inflation-employment gap tradeoff depends heavily on the speed at which labor force participation adjusts to the unemployment rate, as well as on the Phillips Curve slope.

Rather than take the labor force participation gap as an exogenous initial condition as in the static model, we generate initial conditions 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 two years. In all of our simulations, the shock consists of a persistent change in the discount factor ε_{dt} that causes the potential real interest rate r_t^{pot} to fall sharply.

The dash-dotted blue lines in Figure 9 show the effects of the demand shock under a Taylor rule that aims to fully stabilize the unemployment rate at its long-run rate ($u_{rt} = u_{rt}^{pot}$); thus, the Taylor rule given by equation (17) 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. Because the zero bound does bind in our simulation, the unemployment rate u_{rt} rises below u_{rt}^{pot} (panel 1), the participation rate falls persistently (panel 2), and inflation declines below target (panel 4). The policy of targeting the unemployment rate is consistent with a smooth convergence of u_{rt} to u_{rt}^{pot} as the shock wears off and the potential real interest rate (panel 6) rises. Even so, while the unemployment rate gap converges within

10 quarters, the effects on labor force participation are much more long-lived. The highly persistent participation gap exerts downward pressure on inflation, so that inflation remains below target long after the employment rate gap has closed.

Figure 9 also shows the implications of alternative policy rules aimed at strictly stabilizing inflation – the dashed red lines – or the overall employment gap $n_t - n_t^{pot}$ – the solid black lines (achieved by setting γ_π or γ_n to arbitrarily large values, respectively). Consistent with our discussion of the static model, the rule that attempts to stabilize inflation implies an overshooting of the unemployment rate gap, and induces a somewhat faster recovery in both the participation rate gap and in inflation than the rule focused on the unemployment rate. Finally, the rule that attempts to stabilize the overall employment gap n_t induces a persistent overshooting of *both* the unemployment rate (i.e., $u_{rt} < u_{rt}^{pot}$) and inflation as the economy recovers. Clearly, this policy rule also mitigates the size of the initial rise in the unemployment rate and fall in inflation relative to the other policy rules, and promotes a much faster recovery in the participation gap. As we discuss below, the overshooting of the unemployment gap and inflation is qualitatively similar to that under an optimal full commitment policy rule, even though the non-inertial form of the Taylor rule does not embody any sort of commitment to take account of past outcomes.

We next examine some key structural factors influencing the nature of the tradeoff between stabilizing the overall employment gap and inflation in our model. Figure 10 highlights the role of the speed of adjustment of the labor force participation rate to unemployment fluctuations by considering a smaller autoregressive parameter of $\mu_1 = .75$ (dashed red lines) in addition to the baseline of $\mu_1 = 0.95$. In each case, the monetary rule aims to fully stabilize the overall employment gap. While the unemployment rate still overshoots in the case of more rapid adjustment of participation, the small magnitude and transient nature of the overshooting induces inflation to converge smoothly to baseline from below. Thus, the simulation emphasizes very little inflation-overall employment gap tradeoff if labor force participation adjusts quickly to unemployment.

Figure 11 considers how the tradeoff is impacted by a steeper slope of the Phillips Curve. This is achieved by setting the contract duration parameter $\xi_p = .96$ rather than 0.985 as in our baseline. With a steeper Phillips Curve, a policy of strictly stabilizing the overall employment gap implies a sizeable rise in inflation.

We next assess how an optimal “full commitment” policy rule derived from quadratic loss function compares to the Taylor rule that responds to overall employment gap. As noted above, the Taylor rule responses share certain key characteristics of the optimal responses: under both rules, the unemployment rate gap and inflation persistently overshoot target. Even so, the optimal policy provides considerably more stimulus by “promising” an eventual overshooting of the overall employment gap, so that n_t rises above n_t^{pot} . Moreover, the participation rate gap also eventually overshoots as the economy recovers. This behavior clearly contrasts with the non-inertial Taylor rule, which implies smooth convergence of both the overall employment gap and participation gap from below.

5. Conclusions

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), and consequentially, that monetary rules developed for the Great Moderation period may have to be adapted to account for broader measures of slack. But there is much more work to be done. Clearly, there are many sources of uncertainty that may affect the tradeoff between employment gap and inflation stabilization, and that are very relevant in practice for crafting an appropriate strategy.

First, our model does not take account of uncertainty about the natural rate of unemployment, which is likely to be especially large in the wake of a deep recession (due to e.g., a deterioration in the matching function). An extensive literature has emphasized the inflation risk of attempting to target a natural rate that is too low, including research by Orphanides (2002) linking the high inflation of the 1970s to natural rate misperceptions. Second, there

is little evidence – theoretical or empirical – on how a cyclical participation gap is likely to affect the wage demands of workers, and thus firm marginal costs. In the plausible case in which non-participants have a negligible impact on wage demands in the market sector, the upward pressure on inflation associated with overshooting would be greater. Third, while the surprising resilience of inflation during the global recession suggests a very flat Phillips Curve, it is quite possible that the Phillips Curve is more upward-sloping, which would again imply more inflationary pressure associated with a strategy of overshooting unemployment. Fourth, it is possible that the labor force participation gap is smaller than we have estimated – or reacts more quickly to unemployment – either of which would again imply more inflationary risk of overshooting.


References

- Adam, Klaus, and Roberto M. Billi (2008). "Monetary Conservatism and Fiscal Policy." *Journal of Monetary Economics* 55(8), 1376-1388.
- Aaronson, D., Davis, J., Hu, L. 2012. "Explaining the Decline in the U.S. Labor Force Participation Rate." Chicago Fed Letter #296, Federal Reserve Bank of Chicago.
- Aaronson, D., Park, K., Sullivan, D. 2006, "The Decline in Teen Labor Force Participation," *Economic Perspectives*, Federal Reserve Bank of Chicago, 30:2-18.
- Aaronson, S., Fallick, B., Figura, A., Pingle, J., Wascher, W. 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, D., Duggan, M. 2003. "The Rise in the Disability Rolls and the Decline in Unemployment." *Quarterly Journal of Economics*, 157-205.
- Autor, D., Duggan, M. 2006. "The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding." *Journal of Economic Perspectives*, 20:71-96.
- Autor, D. 2010. "The Polarization of Job Opportunities in the U.S. Labor Market." Washington DC: Center for American Progress.
- Bound, J., Waidmann, T. 2002. "Accounting for Recent Declines in Employment Rates among Working-Aged Men and Women with Disabilities." *Journal of Human Resources*, 37:231-250.
- Clark, K., Summers, L. 1982. "Labor Force Participation: Timing and Persistence." *Review of Economic Studies*, 49:825-844.
- Congressional Budget Office. 2013. "The Budget and Economic Outlook: Fiscal Years 2013-2022." February.
- Daly, M., Elias, E., Hobijn, B., Jorda, O. 2012. "Will the Jobless Rate Drop Take a Break?" Economic Letter 2012-37, Federal Reserve Bank of San Francisco.
- Daly, M., Hobijn, B., Şahin, A., Valletta, R. 2012. "A Rising Natural Rate of Unemployment: Transitory or Permanent?" *Journal of Economic Perspectives*.
- DiCecio, Riccardo, Kristie M. Engemann, Michael T. Owyang, and Christopher H. Wheeler. 2008. "Changing Trends in the Labor Force: A Survey." Federal Reserve Bank of St. Louis, *Review*, January/February, pp. 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.
- Elsby, M., Hobijn, B., Sahin, A. 2012. "On the Importance of the Participation Margin for Labor Market Fluctuations." Manuscript, University of Edinburgh.

- Fujita, S. 2011. “Effects of Extended Unemployment Insurance Benefits: Evidence from the Monthly CPS.” Federal Reserve Bank of Philadelphia Working Paper 10-35/R.
- Goldin, C. 2006. “The Quiet Revolution That Transformed Women’s Employment, Education, and Family.” *American Economic Review*, 96:1-21.
- Gruber, J. 2000. “Disability Insurance Benefits and Labor Supply.” *Journal of Political Economy* 58:1162–1183.
- Heckman, J., Killingsworth, M. 1986. “Female Labor Supply: A Survey.” In: O. Ashenfelter and R. Layard, eds., *Handbook of Labor Economics* 1, 103-204. Amsterdam: Elsevier Science Publishers.
- Heim, B. 2007. “The Incredible Shrinking Elasticities: Married Female Labor Supply 1978-2002.” *Journal of Human Resources*, 42:881-918.
- Hotchkiss, J. 2005. “Changes in Behavioral and Characteristic Determination of Female Labor Force Participation, 1975-2005.” *Economic Review*, Federal Reserve Bank of Atlanta, 1-13.
- Hotchkiss, J., Pitts, M., Rios-Avila, F. 2012. “A Closer Look at Non-Participants during and after the Great Recession.” Federal Reserve Bank of Atlanta Working Paper No. 2012-10.
- Hotchkiss, J., Rios-Avila, F. 2013. “Identifying Factors behind the Decline in the Labor Force Participation Rate.” Manuscript, Federal Reserve Bank of Atlanta.
- Juhn, C. 1992. “Decline of Male Labor Market Participation: The Role of Declining Market Opportunities.” *Quarterly Journal of Economics*, 107:79-121.
- Kudlyak, M., Schwartzman, F. 2012. “Accounting for Unemployment in the Great Recession: Nonparticipation Matters.” Federal Reserve Bank of Richmond Working Paper No. 12-04.
- Leonard, J. 1979. “The Social Security Disability Program and Labor Force Participation.” National Bureau for Economic Research Working Paper 392.
- Macunovich, D. 2010. “Reversals in the Patterns of Women’s Labor Supply in the United States, 1977–2009.” *Monthly Labor Review* 133: 16–36.
- Mincer, J. 1966. “Labor Force Participation and Unemployment: A Review of Recent Evidence.” In: R. A. Gordon and M. S. Gordon, eds., *Prosperity and Unemployment*. 73-112.
- Moffitt, R. 2012. “The Reversal of the Employment-Population Ratio in the 2000s: Facts and Explanations.” *Brookings Papers on Economic Activity*, 201-250.
- Oaxaca, R. 1973. “Male-Female Wage Differentials in Urban Labor Markets.” *International Economic Review* 14: 693–709.

- Orphanides, A. 2002. "Monetary Policy Rules and the Great Inflation." *American Economic Review* 92, 115-120.
- Pencavel, J. 1986. "Labor Supply of Men: A Survey." In: O. Ashenfelter and R. Layard, eds., *Handbook of Labor Economics* 1, 3-101. Amsterdam: Elsevier Science Publishers.
- Sherk, J. 2012. "Not Looking for Work: Why Labor Force Participation Has Fallen During the Recession." *Backgrounder* #2722, Heritage Foundation Center for Data Analysis.
- Shimer, Robert. 2008. "The Probability of Finding a Job." *American Economic Review*, 98:268-273.
- Smets, Frank and Raf Wouters (2007). "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97(3), 586-606.
- Smith, C. 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, J., Williams, J. 2010. "Simple and Robust Rules for Monetary Policy." Federal Reserve Bank of San Francisco Working Paper 2010-10.
- Toossi, M. 2007. "Labor Force Projections to 2016: More Workers in their Golden Years." *Monthly Labor Review*, November 2007, 33-52.
- Veracierto, M. 2008. "On the Cyclical Behavior of Employment, Unemployment and Labor Force Participation." *Journal of Monetary Economics*, 55:1143-1157.
- Van Zandweghe, W. 2012. "Interpreting the Recent Decline in Labor Force Participation." *Economic Review*, Federal Reserve Bank of Kansas City, 5-34.
- Woodford, Michael (2003). *Interest and Prices*. Princeton, N.J.: Princeton University Press.

Figure 7: Monetary Policy Tightening
Transient Tightening



Persistent Tightening




Figure 8: Unemployment Gap and Marginal Cost




Figure 9: Adverse Demand Shock under Different Monetary Rules

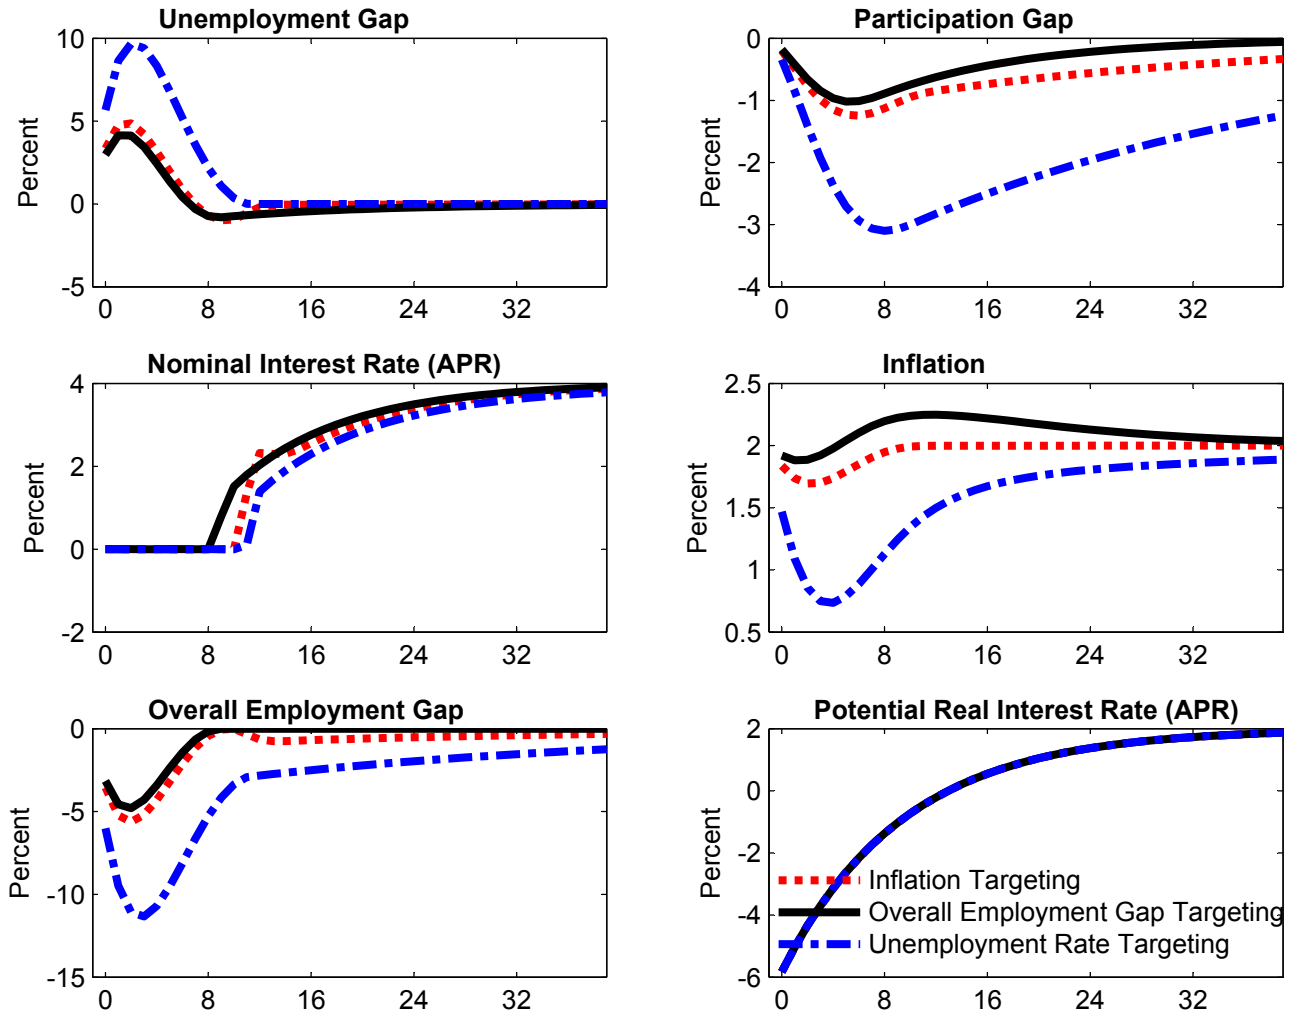


Figure 10: Adverse Demand Shock w/Different Speed of Participation Response

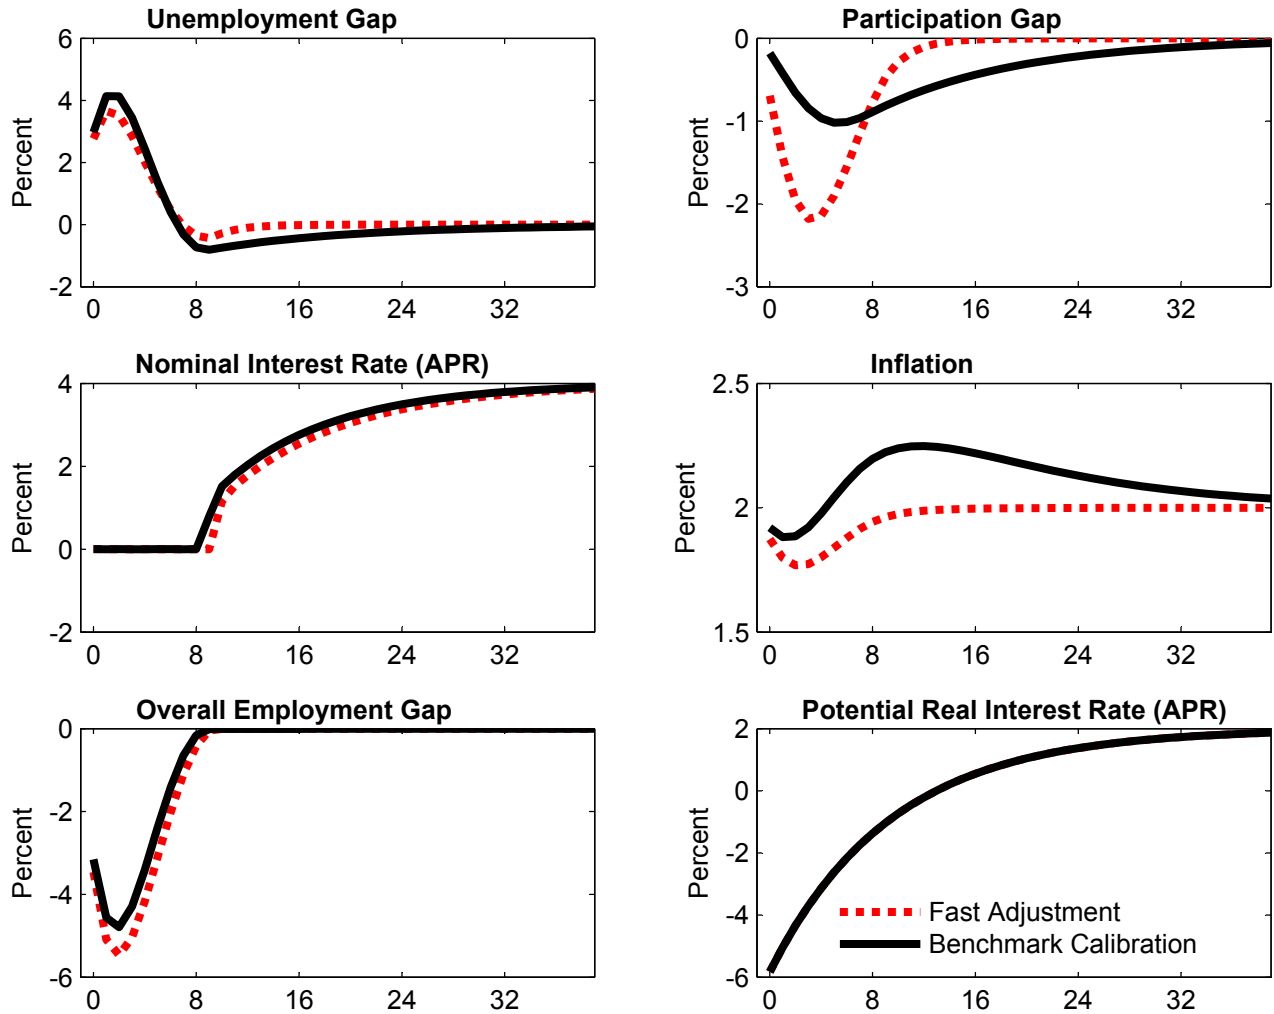
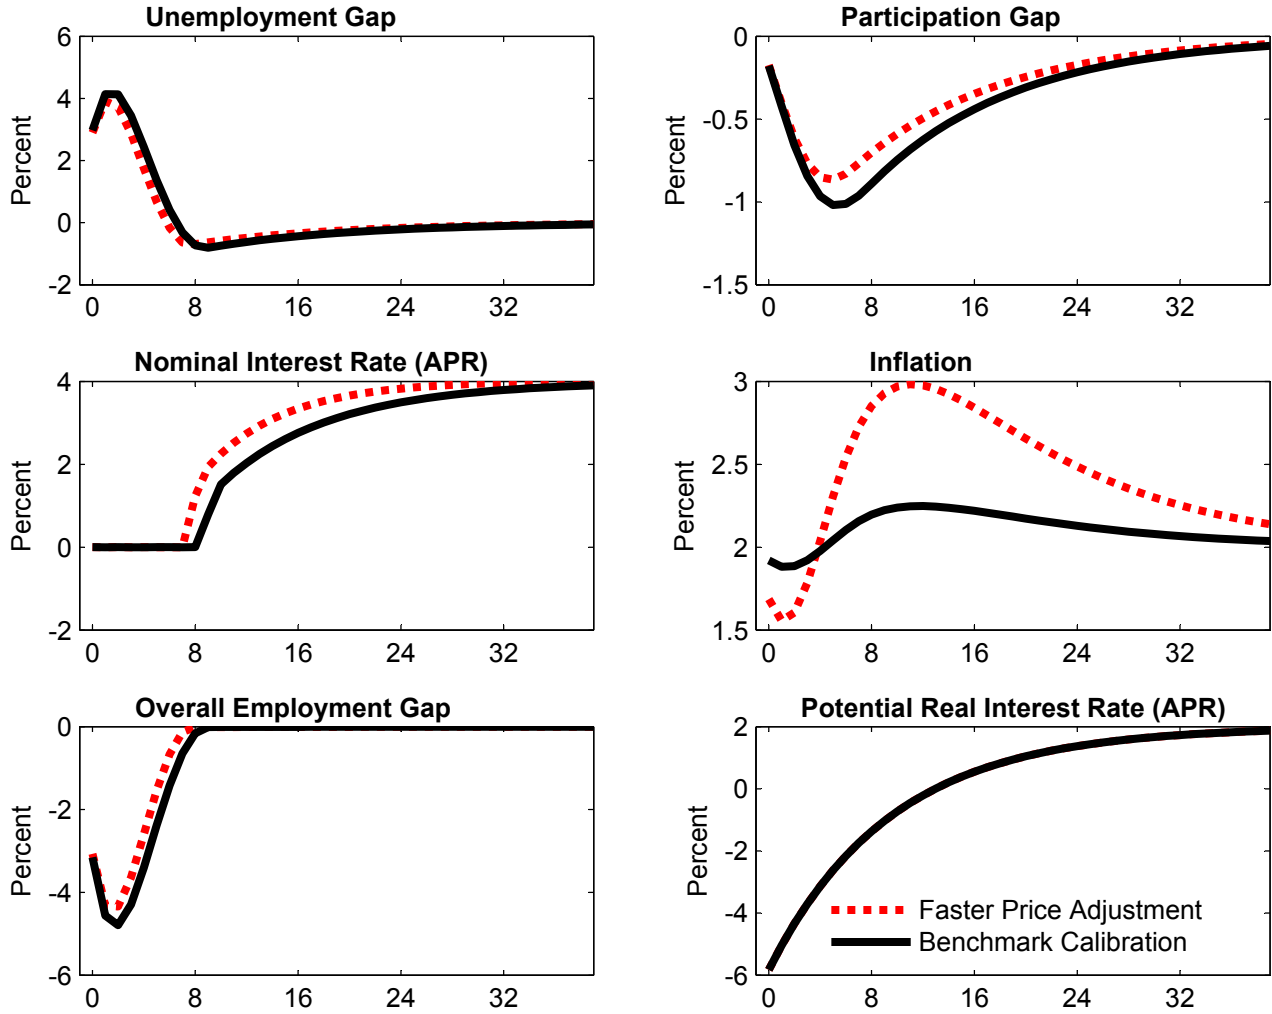



Figure 11: Adverse Demand Shock w/Different Phillips Curve Slopes

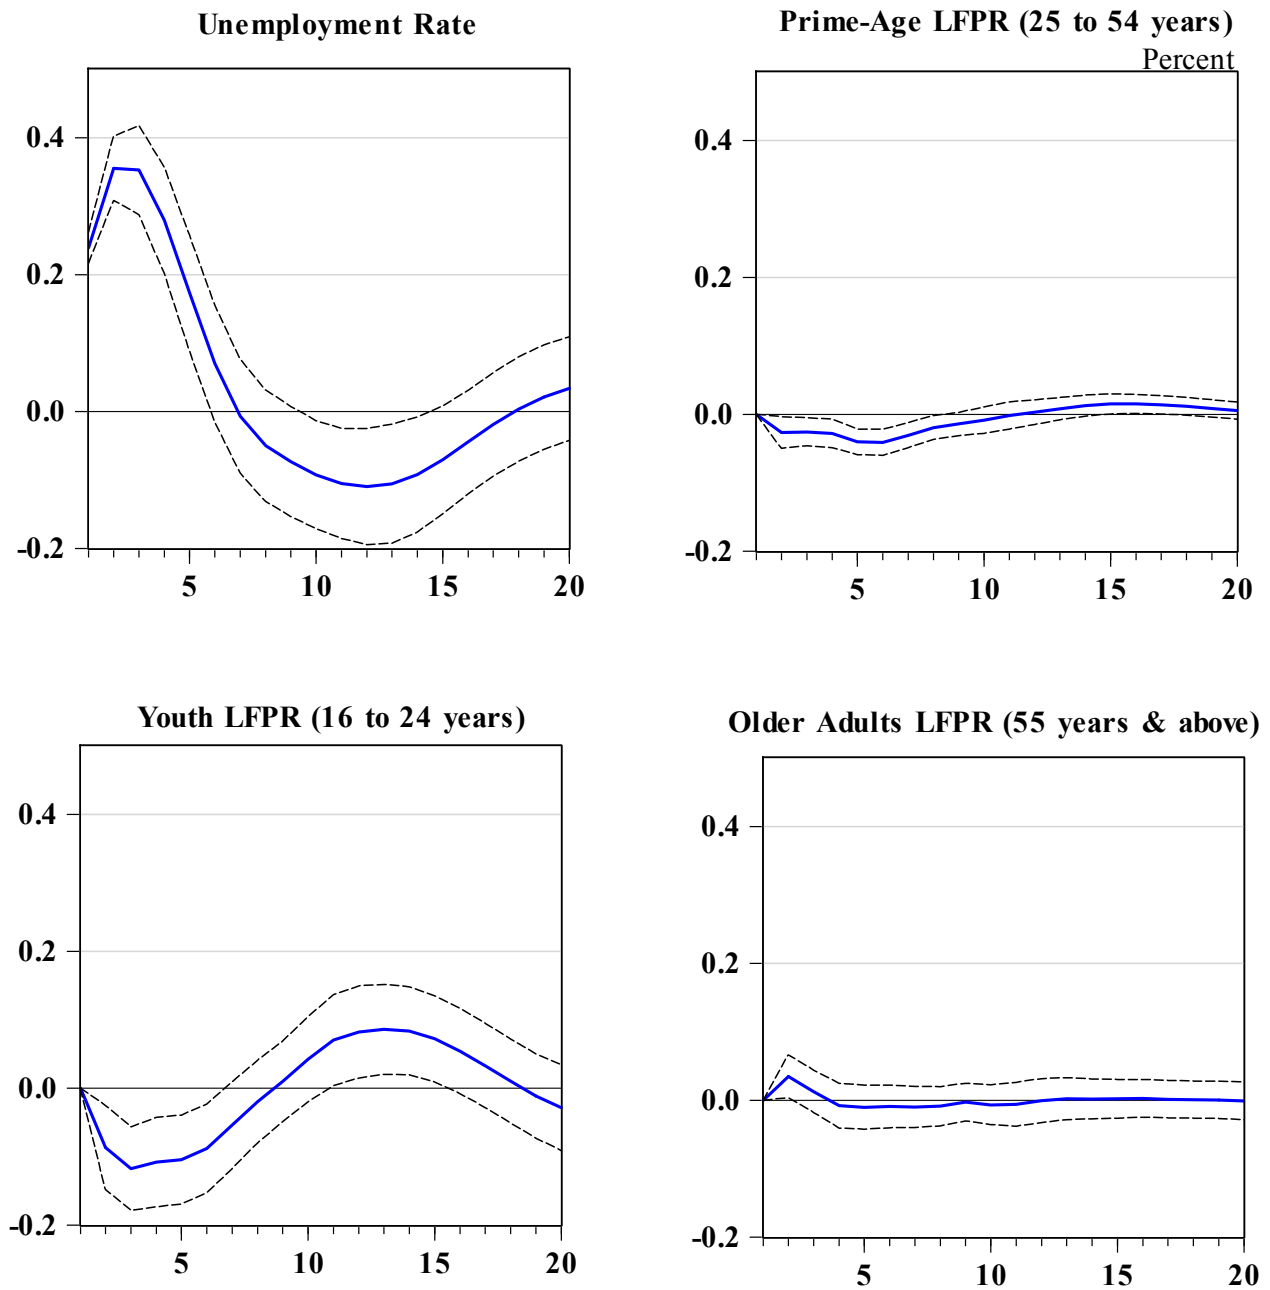


Appendix Figure A1: The Demographic Composition of the Labor Force, 1948 to 2012



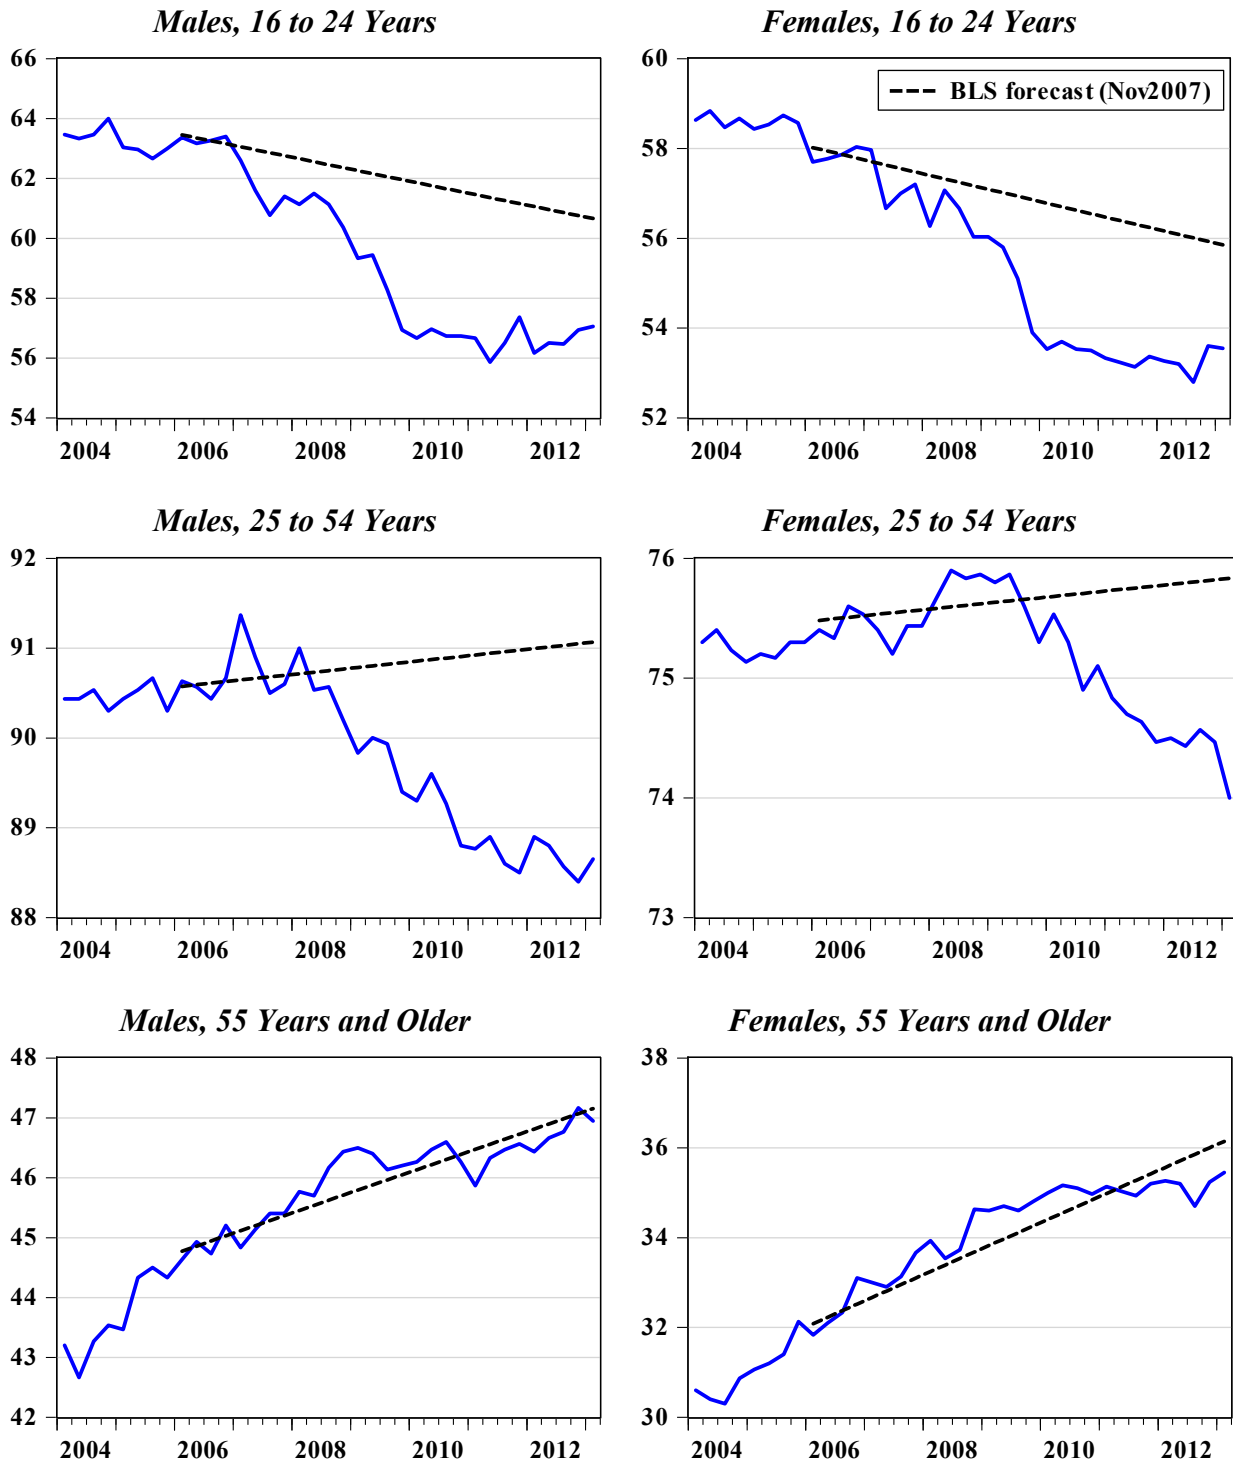
Note: This figure depicts the demographic composition of the labor force (in percent) from 1948 to 2012. The solid area denotes the share of prime-age males (25 to 54 years old), and the horizontal-dashed area denotes the share of prime-aged females (25 to 54 years). The diagonal-dashed and vertical-dashed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

**Appendix Figure A2: The Cyclical Dynamics of the LFPR During “Normal Times”
by Age Group**




Note: This figure depicts the impulse responses from a VAR estimated using quarterly data on the unemployment rate and the labor force participation rates for three demographic groups (prime-age adults, youths, and older adults) over the period 1948:Q1 to 2007:Q4. In each panel, the solid line denotes the response of the specified variable to an orthogonalized unemployment rate shock of one standard deviation, and the dashed lines depict 95 percent confidence bands based on 10,000 Monte Carlo replications.

Appendix Figure A3: Labor Force Participation Rates by Gender and Age



Note: 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*. *Source:* Bureau of Labor Statistics.

**Appendix Figure A4: Data from Larger States (excluding Nevada)
Regarding Unemployment and Labor Force Participation of Prime-Age Adults**



Note: The vertical axis refers to the change in the labor force participation rate for prime-age adults (25 to 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; i.e., 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. *Source:* Bureau of Labor Statistics.