Working Paper Series
Congressional Budget Office
Washington, D.C.

A Markov-Switching Model of the Unemployment Rate

Michael McGrane
Congressional Budget Office
michael.mcgrane@cbo.gov

Working Paper 2022-05

March 2022

To enhance the transparency of the work of the Congressional Budget Office and to encourage external review of that work, CBO's working paper series includes papers that provide technical descriptions of official CBO analyses as well as papers that represent independent research by CBO analysts. Papers in this series are available at http://go.usa.gov/xUzd7.

This paper has not been subject to CBO's regular review and editing process. The views expressed here should not be interpreted as CBO's.

I thank James Hamilton, Nellie Liang, Chris Nekarda, Nicolas Petrosky-Nadeau, and David Wilcox for their invaluable comments and suggestions. I am grateful to members of CBO's Panel of Economic Advisers, especially Alan Auerbach, Emi Nakamura, Jonathan Parker, Valerie Ramey, James Stock, and Kevin Warsh, for their feedback. This work has benefited greatly from discussions with Chris Adams, Gloria Chen, Meredith Decker, Mark Doms, Justin Falk, Mike Falkenheim, Sebastien Gay, Jeffrey Kling, Jeff Perry, Hark Yoo, and other members of CBO.

www.cbo.gov/publication/57582

A Markov-Switching Model of the Unemployment Rate

Michael McGrane

Abstract

The unemployment rate has asymmetric dynamics: It increases rapidly in recessions and falls gradually in expansions. The Congressional Budget Office developed a Markov-switching model to help incorporate these dynamics into macroeconomic projections and cost estimates that require simulations of the national unemployment rate. The model produces simulations that match observed asymmetric business-cycle dynamics at a rate consistent with historical data. I also show that indirect duration dependence, in which transition probabilities are a function of the unemployment gap, creates significant distortions for statistical tests of duration dependence in the business cycle. I present evidence that the benchmark Markov-switching model produces forecasts superior to those of a simpler model with constant transition probabilities, in addition to the linear version of the model. Finally, I make adjustments to the model to account for the unique split between permanently separated unemployment and temporarily separated unemployment in the pandemic recession and recovery.

JEL Classification: C22, C24, C53, E24, E32

1 Introduction

The unemployment rate tends to rise much faster in recessions than it falls in expansions (see Figure 1). However, any linear model with independent and identically distributed shocks implies that the modeled time series will not exhibit this behavior. Although linear models are well suited for producing mean and modal forecasts, they are less useful for exploring alternative paths around these central tendencies. Moreover, macroeconomic simulations using linear models produce recessions much shorter and shallower than those observed during the postwar period.[1]

Many policies that the Congressional Budget Office analyzes condition spending on the path of the unemployment rate. For example, extended benefits for federal/state unemployment insurance may activate if a state's insured unemployment rate equals or exceeds 5 percent and equals or exceeds 120 percent of the average rate over the same period in the last two years.[2] If changes in the unemployment rate are symmetric, then a standard linear model with independent and identically distributed shocks should produce simulations of the unemployment rate that activate these triggers at a rate close to the historical frequency with which these conditions have been met. If, instead, the changes in the unemployment rate are asymmetric, these simulations may not provide a consistent estimate of the frequency with which these conditions will be met over a given forecast horizon.[3] Therefore, producing simulations that are consistent with observed business-cycle dynamics is critical to estimating the cost of programs that condition spending on the path of the unemployment rate.

To capture asymmetric dynamics of the business cycle, CBO developed a Markov-switching model of the unemployment rate. Hamilton (1989) pioneered the use of Markov-switching

[1]Gonzalez-Astudillo and Vilan (2019) documented this finding in the context of the Federal Reserve Board's large-scale macroeconometric model, FRB/US.

[2]For this program, CBO uses complementary models to produce simulations of the national insured unemployment rate and state insured unemployment rates, using simulations of the national unemployment rate as inputs.

[3]The level of the unemployment rate when these triggers are activated also affects the costs of programs like these, which is one reason why CBO uses simulations of these variables instead of just focusing on probabilities.

Figure 1. Quarterly Change in Unemployment Rate

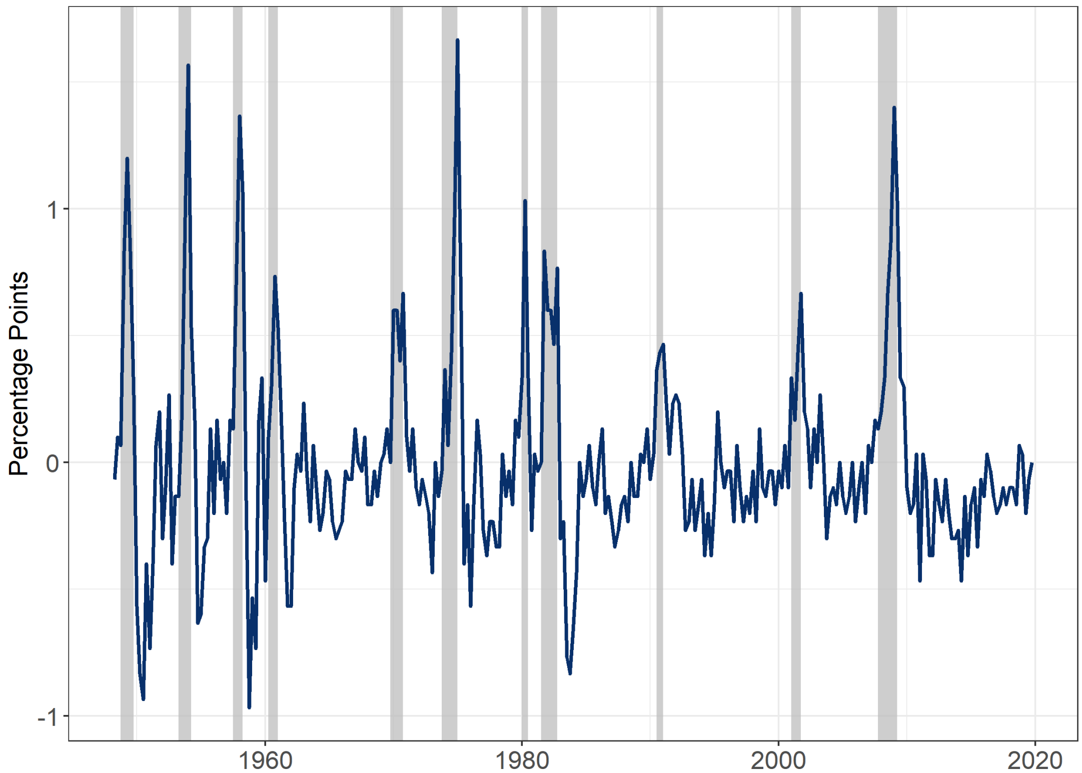

Note: Gray bars denote periods identified as recessions by the National Bureau of Economic Research. The sample is truncated to include only prepandemic data to better show asymmetric dynamics.

models in time-series econometrics by examining gross domestic product (GDP) growth in a model with state-dependent means. Although business cycles drive the behavior of many macroeconomic variables, the unemployment rate represents one of the best indicators of whether the economy is in recession or expansion.[4]

A Markov-switching model allows the economy to be in one of several states at any given time, unlike linear models, which assume a single state. Each state in the Markov-switching model has its own set of parameters governing the data-generating process, and the economy transitions between states with probabilities governed by a Markov chain. Thus, the

[4]Detailed data on unemployment rates across demographic groups also provide useful information about the state of the business cycle, though the data samples for individual groups can be somewhat limited.

economy can transition between expansion states and recession states. Moreover, these states can be estimated endogenously, without prior information supplied by the econometrician beyond the data used to fit the model. That is, the estimation procedure attaches weights to each state across time in a way that maximizes the likelihood of the sample process, and these weights coincide closely with the National Bureau of Economic Research's (NBER's) designations of recession and expansion.

I estimate the model using CBO's historical estimate of the unemployment gap, which effectively removes a low-frequency trend from the observed unemployment rate. I also allow the probability of transitioning between states to be a function of the unemployment gap, instead of keeping these probabilities constant over time. The estimated model produces simulations of the unemployment rate that rise sharply in recessions and decline slowly in expansions, matching the historical data.[5] Moreover, these results are robust to my choice of estimating the model with the unemployment gap (rather than the unemployment rate) and to my assumption that the probability of transitioning between states is a function of the unemployment gap (instead of constant).

I show that the Markov-switching model produces simulations that can be defined as recessions at a rate consistent with the historical data, whereas a simple linear version of the model cannot. I also show that indirect duration dependence, in which transition probabilities are a function of the unemployment gap, creates significant distortions for statistical tests of duration dependence in the business cycle. In the estimated model, the probability of transitioning from expansion to recession increases as the unemployment rate falls. Because the unemployment rate tends to decline over time in expansions, the probability of transitioning to recession will be indirectly related to the duration of the expansion. This feature, along with small sample sizes of postwar expansions and recessions, reduces the power of statistical tests of duration dependence. To demonstrate this, I produce many samples of durations of expansions and recessions using the Markov-switching model and run a statistical test for duration dependence on each sample. I show that the statistical test has almost

[5] Hall and Kudlyak (2020) demonstrated that declines in the unemployment rate during expansions are remarkably consistent across business cycles in the postwar data.

no power to identify duration dependence when the sample size is small and the effect of duration dependence is indirect.

Although the Markov-switching model is not used to produce CBO's baseline projections of the unemployment rate, I provide evidence that the model produces forecasts superior to those of a simple linear version of the model.[6] Specifically, I use the Giacomini-White (2006) test (hereafter, *GW test*) to compare the forecasts produced by the Markov-switching model with those of a simple AR(2) model of the unemployment gap, in addition to a simpler version of the Markov-switching model that uses constant transition probabilities (CTPs). I find that the benchmark Markov-switching model produces smaller out-of-sample root mean square errors (RMSEs) than either alternative model beyond the 10-quarter forecast horizon, with varying degrees of statistical significance. To evaluate the ability of the model to forecast periods that may be defined as recessions, I conduct a similar test comparing the frequencies with which the benchmark model produces simulations that trigger recession thresholds with the frequencies produced by the simpler models. Again, the benchmark model outperforms the alternative models in forecasting periods that may be defined as recessions at statistically significant levels. Additionally, I use the Giacomini-Rossi (2010) test (hereafter, *GR test*) to evaluate whether the relative performance of the models has changed over time. I find that the relative forecasting performance of the benchmark model increased significantly before the 2007–2009 recession, as the model tended to produce higher medium-term forecasts of the unemployment rate than the alternative models.

Despite the ability of the Markov-switching model to capture the dynamics of the unemployment rate in recession and expansion, the recession related to the COVID-19 pandemic presents unique challenges for time-series models. To account for the dramatic changes in the labor market since the start of that recession, I separately model the unemployment rate for workers who have been temporarily separated from their employer and the unemployment rate for workers who have experienced a permanent separation from their previous employer when creating projections and simulations over the near term. This allows me to more accu-

[6]For a thorough description of how CBO constructs its macroeconomic forecasts, see Arnold (2018).

rately characterize the duration of unemployment for newly unemployed workers. Moreover, the modeling framework involves the use of the Markov-switching model of the unemployment gap for creating unemployment-rate simulations when the split between temporarily and permanently separated unemployment normalizes. This procedure produces simulations that are consistent both with the particular dynamics observed since the start of the pandemic recession over the near term and with past business-cycle dynamics over the longer term.

2 Evidence of Asymmetry and Potential Sources

2.1 Evidence of Asymmetry

Business-cycle asymmetry was noted by early business-cycle researchers, including Keynes (1936) and Burns and Mitchell (1946). However, asymmetry was not examined empirically until the early 1980s, when the research program of documenting stylized facts of business cycles began.[7] Using a second-order Markov process, Neftçi (1984) proposed a statistical test for the hypothesis that a variable rises faster than it falls and found evidence of this property in the unemployment rate.[8] DeLong and Summers (1986) reported similar evidence that the unemployment rate rises faster than it falls. However, Falk (1986), using Neftçi's test, found that this asymmetry was not present in the growth of real output—that is, real output does not fall faster than it rises. Sichel (1993) proposed an alternative test for asymmetry that involved estimating the skewness of the rate of change of a variable under consideration. If the skewness in the change of a variable is positive, then the variable will tend to rise faster than it falls. Confirming earlier work, Sichel found that the unemployment rate exhibited this asymmetry, whereas growth in real gross national product (GNP) did not. Later work by McKay and Reis (2008) also found that contractions in employment are briefer and more

[7] In this paper, I focus on the asymmetry of the unemployment rate's rising faster than it falls, though there is a large literature on other types of asymmetries.

[8] Sichel (1989) discovered a mistake in the empirical work that made the results appear statistically significant when in fact they were not. However, Rothman (1991) showed that the results are indeed statistically significant if a first-order Markov process is used instead.

violent than expansions, while equal brevity and violence for expansions and contractions in output could not be rejected. Thus, although there is little evidence for asymmetry in real output growth, it is clear that the unemployment rate tends to rise faster than it falls.

2.2 Potential Sources of Asymmetry

Two mechanisms have been proposed that may help explain this business-cycle asymmetry. The first involves labor productivity shocks in variants of the Diamond-Mortensen-Pissarides (DMP) search model of the labor market (see, for example, Diamond, 1982; Mortensen and Pissarides, 1994). In these models, unemployed workers search for jobs and employers search for workers to hire, and the two are matched according to some probability. A negative productivity shock results in firms' laying off workers, increasing the unemployment rate. The second mechanism that may explain labor market asymmetry involves downward nominal wage rigidity. In frictionless models of the labor market, workers supply labor and firms demand labor on the basis of the wage rate, with market clearing resulting in an optimal allocation of resources. However, when nominal rigidities are introduced, firms and workers may not be able to supply or demand labor at the socially optimal level, resulting in suboptimal allocations of resources. Either of these mechanisms (or some combination of the two) may account for asymmetric labor market dynamics.

In a DMP model with on-the-job search, Pissarides (1994) showed how search by workers who are already employed leads to equilibria in which the unemployment rate may fall more slowly when economic activity improves. For example, following a positive productivity shock, firms respond by creating new job vacancies. Additionally, some workers who are already employed begin searching for work. This entry into the job market creates congestion for unemployed workers who are searching for work and increases the flow of applications to firms with job vacancies. So, although the overall supply of jobs increases, the unemployment rate does not decline as much as would be expected without on-the-job search. Using a variant of the DMP model, Andolfatto (1997) showed that cyclical movements in labor productivity lead to an asymmetric response of job destruction, whereas the rate of job creation is symmetric. In this model, a negative productivity shock is followed by an increased flow of job destruction that is not matched by a subsequent increased flow in job creation. Similarly, den Haan et al. (2000) showed in a dynamic general equilibrium model that endogenous job destruction following a negative productivity shock can significantly amplify the initial effect of the shock as well as the shock's persistence. Additionally, Kohlbrecher et al. (2016) showed that idiosyncratic shocks in a DMP model generate a selection effect at the time of hiring, leading to an asymmetric response of the job-finding rate. Similarly, Ferraro (2018) showed that in a DMP model with heterogeneous labor productivity, shocks to labor productivity can produce endogenous job destruction, search externalities, and fluctuations in average labor productivity, which lead to sharp increases in the unemployment rate at the beginning of recessions.

Alternatively, Abbritti and Fahr (2013) showed that a New Keynesian model with matching frictions and downward nominal wage rigidity can account for the asymmetric fluctuations in employment. In this model, contractions are associated with significant declines in employment while wages adjust slowly. Meanwhile, expansions are associated with rapid adjustments in wages but gradual adjustments of employment.

Dupraz et al. (2019) merged these two mechanisms, embedding downward nominal wage rigidity in a DMP model with heterogeneous labor markets. Although their main finding was that this model embodies a "plucking" model of business cycles, they also showed that this setup generates asymmetric responses of the unemployment rate to productivity shocks. Although I do not take an explicit stand on the source of this asymmetry in this paper, it is likely that some combination of search frictions and downward nominal wage rigidities generates an asymmetric response of the unemployment rate to shocks.

3 Model, Estimation, and Results

In this section, I present the Markov-switching model used to create simulations of the unemployment rate, the estimation procedure, and the estimated coefficients and dynamics after transition. I show that simulations produced by the model can be defined as recessions at a rate consistent with the historical data. Additionally, I show that indirect duration dependence, in which transition probabilities are a function of the unemployment gap, creates significant distortions for statistical tests of duration dependence in the business cycle.

3.1 Model

The unemployment gap is modeled as an AR(2) process in which the gap is a function of a constant and two lagged values of itself.[9] Formally, let $\widetilde{u}_t$ be the unemployment gap (the difference between the unemployment rate u_t and CBO's estimate of the noncyclical rate of unemployment u_t^*) in period t.[10] Additionally, let s_t be the state of the economy in period t, which can be in either expansion, $s_t = e$, or recession, $s_t = r$. For arbitrary $s_t = i$,

$$\widetilde{u}_t = \alpha_i + \beta_i \widetilde{u}_{t-1} + \gamma_i \widetilde{u}_{t-2} + \sigma_i \varepsilon_t \tag{1}$$

where $\varepsilon_t \sim N(0,1)$ and σ_i is the standard deviation of the state-specific error.

The economy transitions between states with probabilities governed by the time-varying transition matrix:

$$P_t = \begin{bmatrix} p_{e,t} & 1 - p_{e,t} \\ 1 - p_{r,t} & p_{r,t} \end{bmatrix}$$

where the element in the ith row and jth column is the probability of transitioning from state i in period $t-1$ to state j in period t. Figure 2 displays a graphical representation of the Markov chain. So, for example, given that the economy was in expansion last period, $s_{t-1} = e$, the probability that it will stay in expansion this period is $p_{e,t}$, and the probability

[9] A lag length of 2 is used for the model because this lag length minimizes both the Akaike information criterion and the Bayes/Schwarz information criterion for a linear AR process of the unemployment gap.

[10] CBO's estimate of the noncyclical rate of unemployment is defined as the rate that arises from all sources other than fluctuations in demand associated with the business cycle. This rate is effectively a low-frequency trend in the unemployment rate that abstracts from business-cycle fluctuations.

Figure 2. Markov Chain

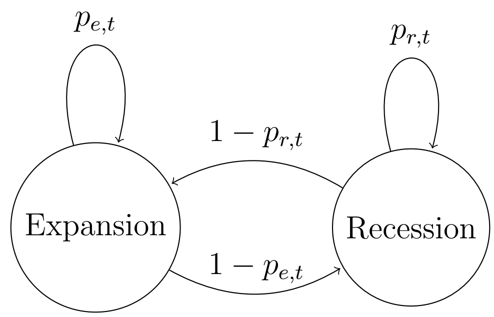

that it will transition into recession is $1 - p_{e,t}$. Similarly, if the economy was in recession last period, $s_{t-1} = r$, the probability that it will stay in recession this period is $p_{r,t}$, and the probability that will transition into expansion is $1 - p_{r,t}$.

The time-varying transition probabilities (TVTPs) are modeled as a function of the lagged unemployment gap and have logistic functional form. The probability of transitioning from state i in period $t-1$ to state i in period t is given by

$$p_{i,t} = \frac{\exp(\delta_i + \zeta_i \widetilde{u}_{t-1})}{1 + \exp(\delta_i + \zeta_i \widetilde{u}_{t-1})} \tag{2}$$

Note that because the transition probability is a function of the unemployment gap, the probability of transitioning between states can vary over the business cycle. With this specification, a value of $\zeta_i > 0$ implies that the probability of staying in state i falls as $\widetilde{u}_{t-1}$ decreases. Thus, if ζ_i exceeds 0 in the expansion state, the probability of staying in the expansion state falls as $\widetilde{u}_{t-1}$ decreases. Alternatively, if ζ_i is less than 0 in the recession state, the probability of staying in the recession state falls as $\widetilde{u}_{t-1}$ increases. In the estimation of the model, no sign restrictions are placed on these parameters.

It is worth noting that the use of the lagged unemployment rate as an explanatory variable for the transition probabilities is not based on some prior belief about the unemployment rate affecting the probability of transitioning between states over the business cycle. Rather, when I constructed simulations using CTPs, many simulations continued to decline beyond minimum historical values. Alternatively, TVTPs that are a function of the unemployment gap exert discipline on how low simulations tend to decline, thus producing simulations that more closely match the historical data.

I do not directly observe which state the economy is in for each time period. Instead, I need to infer its probability of being in each state. These probabilities are formed in the estimation procedure, which simultaneously estimates the parameters of the model.

3.2 Estimation

The model is estimated using an expectation-maximization (EM) algorithm that maximizes the log likelihood of the sample process. This algorithm works by forming state probabilities across history for a given set of parameters and then updating these parameters using the newly formed state probabilities. These updated parameters are then used to form new state probabilities in the next iteration, and so on. This process is repeated until the log likelihood converges. The estimation process is formally presented in Appendix A.

The model is estimated on quarterly data from Quarter 1 (Q1) of 1959 through Quarter 4 (Q4) of 2019. This estimation sample is chosen for two reasons. First, when estimated on data going back to 1949, the model interprets most of the 1950s as a recessionary period. One reason this might be the case is that the variance of the change in the unemployment gap was much higher during the 1950s, which likely leads the model to interpret this time period as recessionary. Moreover, a break-point test of a linear AR(2) model strongly rejects the null hypothesis of no structural break over the entire sample period and suggests that a structural break occurred in Q1 1959.

Initial parameter values are determined through a random search and selection of the parameters associated with the highest initial log likelihood. The EM algorithm is then run, starting with the initial values provided by the random search. I do not impose any sign restrictions on any of the parameters of the model, but rather call the state with the lower constant the "expansion" state and the state with the higher constant the "recession" state.[11]

[11] Moreover, I do not impose any stationary condition in the estimation, but these conditions are satisfied with the resulting estimated parameters.

3.3 Results

Panel A of Table 1 displays the estimated parameters of the model, along with the associated standard errors. These standard errors are formed using White's (1982) covariance matrix, which can be valid even if the probability density used in the maximum likelihood estimation is misspecified. Several features are worth noting. First, the expansion state has a negative constant, whereas the recession state has a positive constant. This indicates that in the absence of shocks, the unemployment gap in the expansion state will tend to fall over time if it is above its state-specific steady state, and the unemployment gap in the recession state will tend to rise over time if it is below its state-specific steady state. Second, the estimated coefficients on the lags of the unemployment gap suggest that the recession state has more pronounced propagation of shocks than the expansion state.[12] Moreover, the variance of the error term is much higher in the recession state than in the expansion state. Finally, the positive coefficient on the lagged unemployment gap for the transition probability in the expansion state suggests that the probability of transitioning from the expansion state to the recession state increases as the unemployment gap falls and decreases as the unemployment gap rises. Alternatively, the negative coefficient on the lagged unemployment gap for the transition probability in the recession state suggests that the probability of transitioning from the recession state to the expansion state increases as the unemployment gap rises and decreases as the unemployment gap falls, though the estimated coefficient is not statistically significant.

To show how the fit of the Markov-switching model compares with that of a standard linear model, Figure 3 compares the fitted residual from the Markov-switching model with the fitted residual from an AR(2) model of the unemployment gap estimated over the same sample. The residual from the Markov-switching model has a slightly smaller standard deviation (0.202) than the linear model (0.234), though this is not surprising, given the larger number of parameters used to fit the Markov-switching model. Although the Markov-switching model

[12] That is, the same shock will have a much more pronounced effect in the recession state than in the expansion state, suggesting some underlying nonlinearity in the effect of macroeconomic shocks to the labor market.

Table 1. Parameter Estimates

A. Baseline Model (Time-Varying Transition Probabilities)

	Dynamics Within State				Transition Probabilities	
	Constant	$\widetilde{u}_{t-1}$	$\widetilde{u}_{t-2}$	σ_i	Constant	$\widetilde{u}_{t-1}$
Expansion State	-0.084^{***} (0.021)	1.163^{***} (0.105)	-0.188^{*} (0.103)	0.126^{***} (0.012)	3.449^{***} (0.468)	0.755^{**} (0.320)
Recession State	0.151^{***} (0.038)	1.581^{***} (0.093)	-0.665^{***} (0.101)	0.334^{***} (0.034)	2.424^{***} (0.485)	-0.086 (0.260)

B. Constant Transition Probabilities

	Constant	$\widetilde{u}_{t-1}$	$\widetilde{u}_{t-2}$	σ_i	Constant
Expansion State	-0.083^{***} (0.024)	1.173^{***} (0.116)	-0.200^{*} (0.111)	0.124^{***} (0.014)	2.953^{***} (0.425)
Recession State	0.152^{***} (0.039)	1.579^{***} (0.094)	-0.658^{***} (0.099)	0.332^{***} (0.036)	2.216^{***} (0.466)

C. Unemployment Rate Instead of Unemployment Gap

	Constant	u_{t-1}	u_{t-2}	σ_i	Constant	u_{t-1}
Expansion State	0.083 (0.057)	1.191^{***} (0.111)	-0.221^{**} (0.108)	0.125^{***} (0.014)	0.413 (0.474)	0.523^{***} (0.118)
Recession State	0.626^{***} (0.157)	1.562^{***} (0.090)	-0.644^{***} (0.092)	0.332^{***} (0.038)	2.293^{***} (0.472)	0.002 (0.005)

Note: The table reports maximum likelihood estimates of model coefficients. Values in parentheses are standard errors calculated using White's (1982) covariance matrix. * $p < 0.1$. ** $p < 0.05$. *** $p < 0.01$.

Figure 3. Fitted Residuals

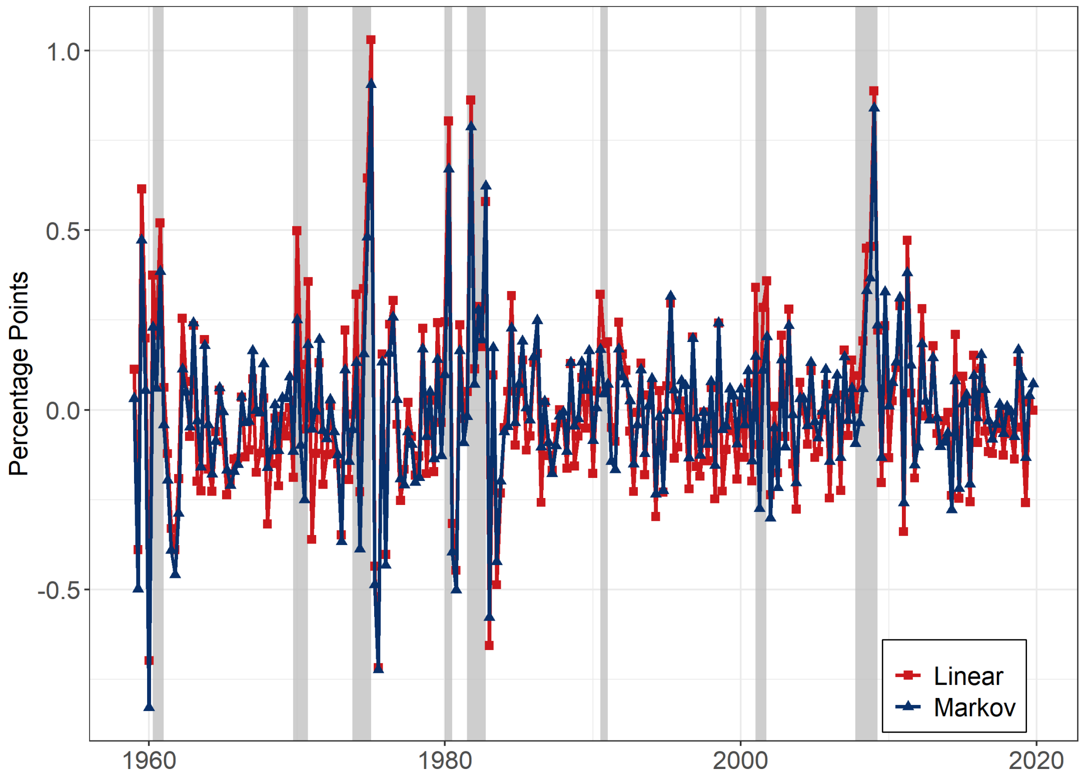

Note: The graph plots the fitted residuals of the Markov-switching model with time-varying transition probabilities and a linear AR(2) model of the unemployment gap. Gray bars denote periods identified as recessions by the National Bureau of Economic Research.

still appears to underpredict the unemployment rate during recessions (the residual appears to be positive, on average, during recessions), the absolute values of the errors tend to be smaller during these periods, reflecting the asymmetric dynamics captured by the model.

To illustrate the dynamics of each state, Figure 4 displays the deviation of the unemployment gap from the model's steady state after transitioning to each state for at least one quarter.[13] Each line represents the average across 1,000 simulations that transitioned to the respective state for at least one quarter, after which the simulations transition endogenously between states. Transition to the recession state is characterized by an increase

[13]The model's steady state is approximated as the average unemployment gap across 10,000 simulations over 80 quarters, allowing for a burn-in period of 40 quarters.

Figure 4. Dynamics of Unemployment Gap After Transition

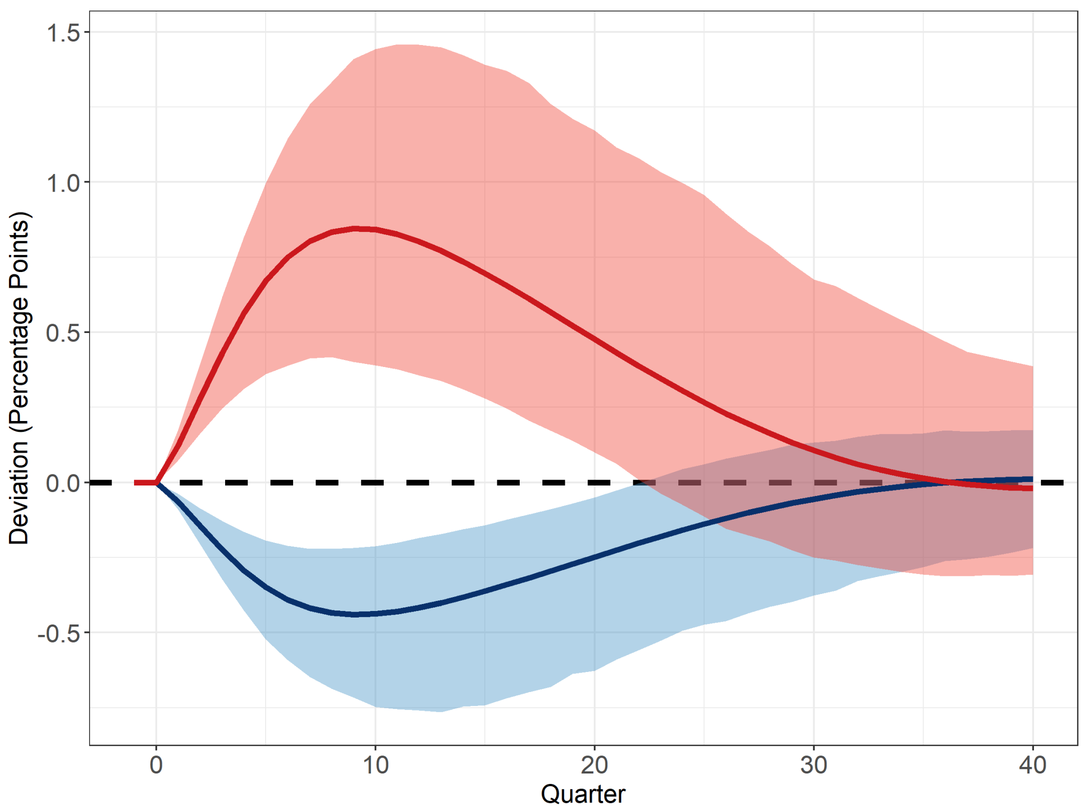

Note: The graph plots the deviation of the unemployment gap from the steady state after transition to each state for at least one quarter. Blue denotes expansion states; red denotes recession states. Solid lines are averages across 1,000 simulations with no shocks using estimated parameters. Shaded regions are 95 percent confidence intervals from parametric bootstrap.

in the unemployment gap of slightly more than 0.8 percentage points over 8 quarters, after which the unemployment gap gradually falls and reaches its initial level after 30 quarters. Alternatively, transition to the expansion state is characterized by a decline in the unemployment gap of 0.4 percentage points over 8 quarters, after which it gradually increases and reaches its initial level after 30 quarters. Although the timing of peaks and troughs and the return to initial levels are roughly the same across states, the unemployment gap in the recession state increases at twice the rate that it falls in the expansion state, matching the historical observation that the unemployment rate increases faster in recessions than it falls in expansions.

Figure 5. Quarterly Transition Probabilities

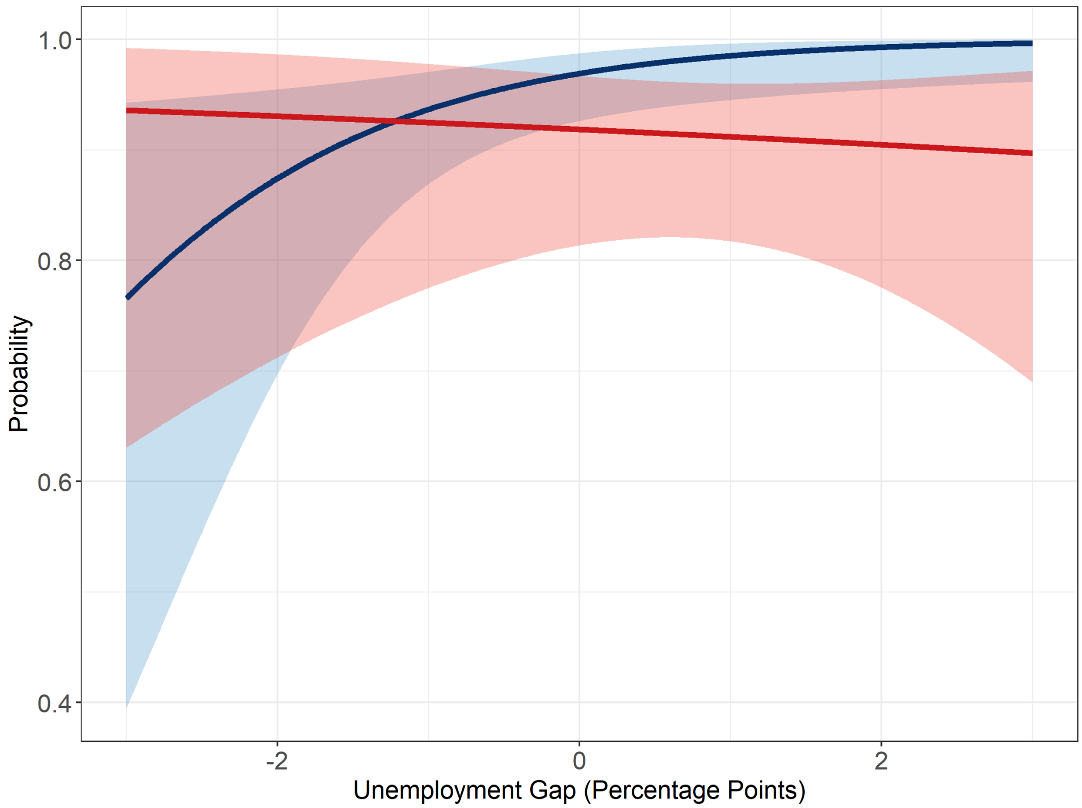

Note: The graph plots the probability of transitioning from state i at time t to state i in time $t+1$ conditional on the unemployment gap. Blue denotes expansion states; red denotes recession states. Solid lines are transition probabilities using point estimates of parameters. Shaded regions are 95 percent confidence intervals from parametric bootstrap.

To provide a better representation of how the transition probabilities change over the business cycle, Figure 5 displays the estimated transition probabilities as a function of the unemployment gap. Conditional on the unemployment gap's being zero, the probability of staying in the expansion state is 0.966, whereas the probability of staying in the recession state is 0.904.[14] Moreover, the probability of transitioning from the expansion state to the recession state increases as the unemployment rate declines. This result is somewhat at odds with earlier literature that found little evidence that expansions die of old age (see,

[14]Conditional on the unemployment gap's being zero, solving for the ergodic state distribution implies that the economy is in the expansion state 74 percent of the time and is in the recession state 26 percent of the time.

Figure 6. Weight on Recession State

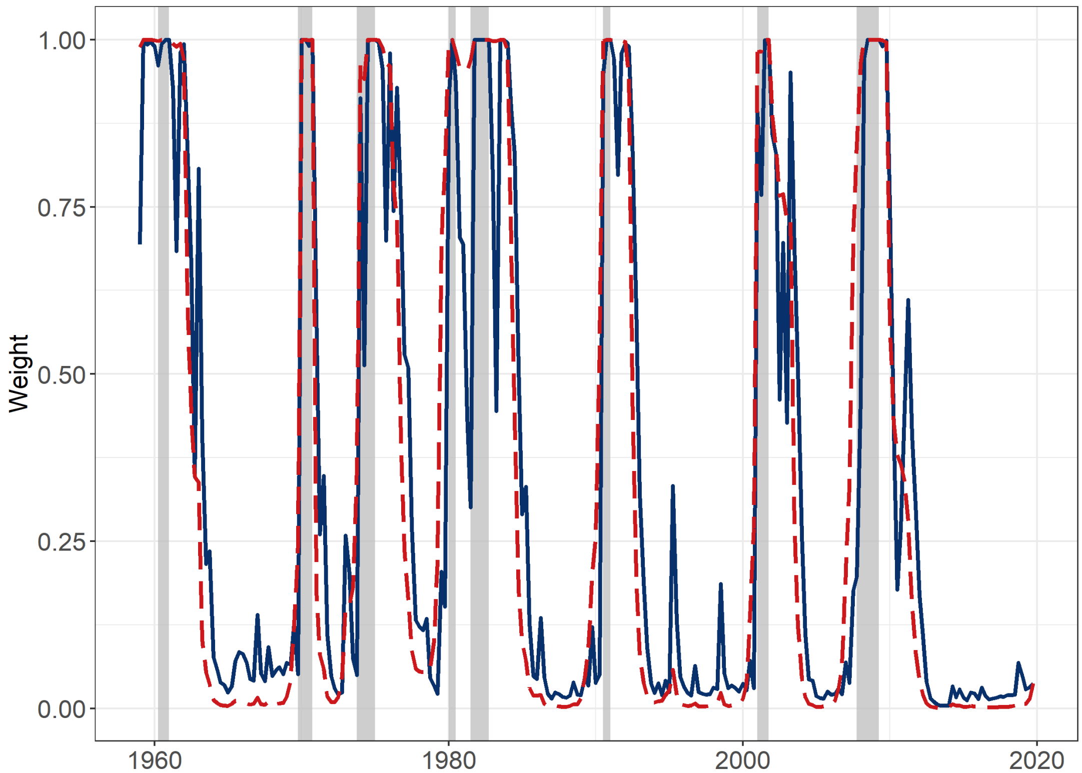

Note: The solid blue line is the filtered probability of being in the recession state at time t; the dashed red line is the smoothed probability of being in the recession state at time t. Gray bars denote periods identified as recessions by the National Bureau of Economic Research.

for example, Diebold and Rudebusch, 1990). Although in this model, the probability of transitioning from expansion to recession is not directly a function of time spent in expansion, that probability increases as the unemployment rate declines, which tends to occur gradually during expansions. The effect of the unemployment gap on transition probabilities has the opposite sign in the recession state, though the magnitude of the effect is much lower (and not statistically significant).

It is also informative to evaluate which periods in history the model identifies as expansions and which it identifies as recessions. To this end, Figure 6 illustrates the weight that the model attaches to the recession state over time. The figure displays two series of estimates that the economy was in the recession state in each quarter. The first, "filtered,"

Figure 7. Probabilities of Recession

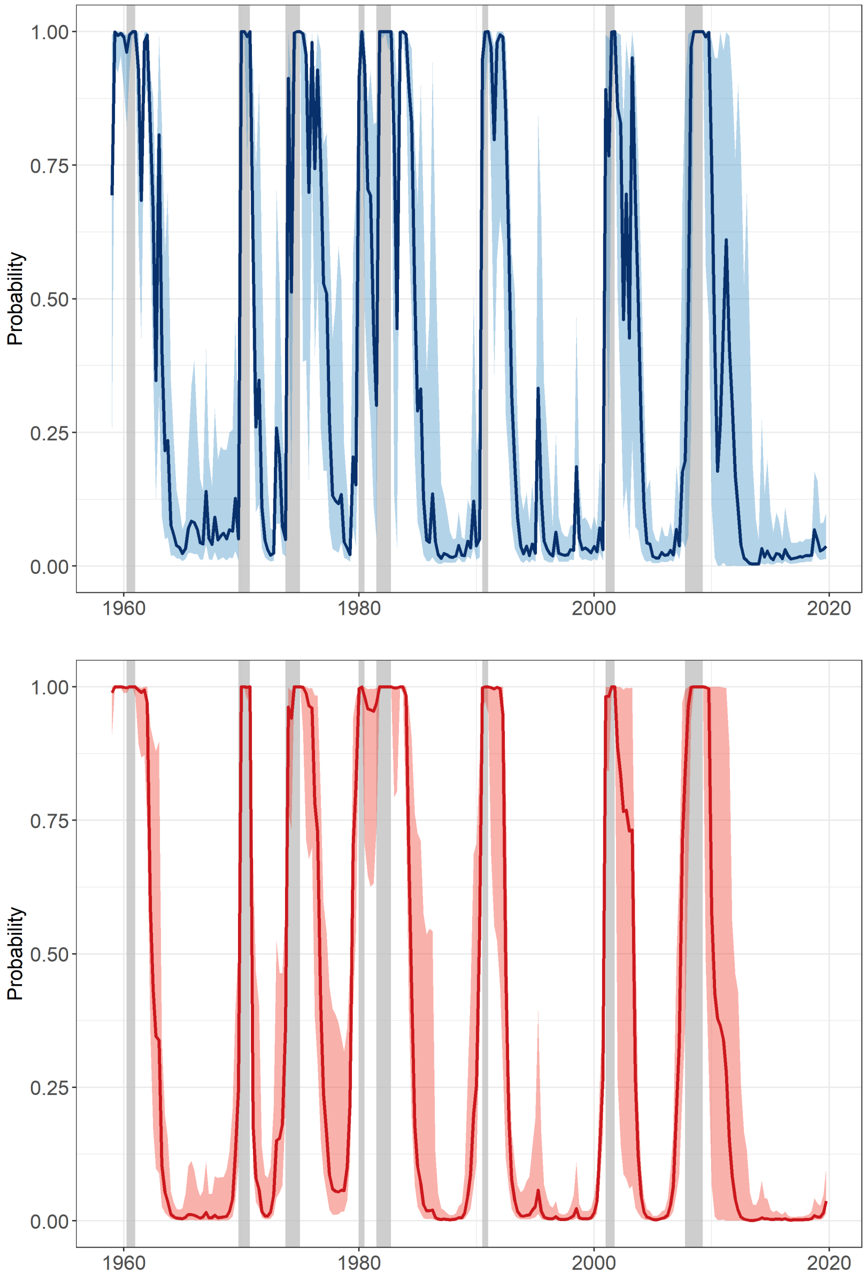

Note: The top panel plots the filtered probability of being in a recession state at time t; the bottom panel plots the smoothed probability of being in a recession state at time t. Shaded regions are 95 percent confidence intervals from parametric bootstrap. Gray bars denote periods identified as recessions by the National Bureau of Economic Research.

is the probability conditional on all observations *before that date*, along with the estimated parameters. The second, "smoothed," is the probability conditional on all observations *over the entire sample*, along with the estimated parameters. That is, the filtered series is purely backward-looking, whereas the smoothed series incorporates future observations when attaching a weight to the recession state at each point in time. Appendix A describes the formulation of these series in more detail. Both measures of the probability of being in the recession state identify periods of history similar to NBER recession dates. The largest discrepancy is that the model tends to continue to attach a high weight to the recession state for several quarters after NBER-dated recessions, reflecting the observation that the unemployment rate tends to peak after the end of recessions.

3.4 Parameter Uncertainty

To demonstrate the effect of parameter uncertainty on the results, Figures 5, 4, and 7 display shaded regions that represent 95 percent confidence intervals. These confidence intervals are constructed using a parametric bootstrap, where parameters are drawn from the multivariate normal distribution with mean equal to the estimated parameters and White's (1982) covariance matrix. For each figure, series are constructed using draws from the bootstrap, with the top of the confidence interval representing series at the 97.5 percent quantile and the bottom of the interval representing series at the 2.5 percent quantile.

The confidence intervals in Figure 5 show that parameter uncertainty has a significant effect on transition probabilities. By contrast, Figure 4 shows that parameter uncertainty has a relatively small effect on the dynamics of the unemployment gap in each state. These findings reflect the relatively small standard errors for the parameters governing the dynamics of the unemployment gap in each state and the higher standard errors for the parameters governing transition probabilities. Finally, Figure 7 displays the effect of parameter uncertainty on the weight that the model attaches to the recession state over time. As expected, parameter uncertainty has less of an effect on the smoothed probability, reflecting the fact that the smoothed probability incorporates observations from the entire sample, whereas the filtered probability incorporate only observations prior to each date.

3.5 Simulations and Dynamics Compared to History

As mentioned earlier, many policies that CBO analyzes condition spending on the path of the unemployment rate. To see what we gain from using the Markov-switching model of the unemployment gap, I produce simulations of the unemployment rate using this model along with simulations produced using a linear AR(2) model of the unemployment gap and then calculate the frequency with which each model produces simulations that may be described as "recessionary."[15]

I produce 100,000 simulations from each model by drawing shocks from the sample of estimated errors. For the Markov-switching model, the probability of drawing shock $\hat{\varepsilon}_t$ in state i is proportional to the probability that the economy was in state i in period t. Moreover, the starting fraction of simulations in each state is set according to the one-step-ahead probabilities in the last estimation period, with the simulations transitioning endogenously between states thereafter. Finally, I allow for a burn-in of 80 quarters before evaluating the simulations to minimize the effect of initial conditions on the estimated dynamics.

To evaluate what fraction of simulations may be defined as "recessionary," I identify recessions within simulations using the Sahm (2019) indicator. This indicator identifies a contraction as occurring when the 3-month moving average of the unemployment rate rises by more than 0.5 percentage points above its minimum value during the previous 12 months. Thus, the Sahm indicator identifies periods in which the unemployment rate is increasing rapidly. Historically, this indicator identifies periods similar to those defined as recessions by the NBER, though contractions identified by the Sahm indicator tend to be longer on average (18.5 months) than those identified by NBER (10.8 months).

Table 2 displays the time spent in recession over the historical sample as identified by the Sahm indicator, along with the share of simulations from each model that are characterized as recessions. The Markov-switching model produces simulations that can be characterized

[15]The linear AR(2) model is estimated over the same sample as the benchmark Markov-switching model.

Table 2. Simulation Results: Proportion of Time Spent in Recessions

History	0.262
Markov-Switching Model	0.266
AR(2) Model	0.368

Note: The table reports the proportion of time spent in recession, as defined by the Sahm indicator. The AR(2) model is estimated using same sample as Markov-switching model. Values reported come from 100,000 simulations over a 40-quarter forecast horizon, following a burn-in of 80 quarters.

as recessions 27 percent of the time, compared with 37 percent of the time for the AR(2) model. By comparison, the Sahm indicator classifies 26 percent of the postwar period as recessions. This suggests that the Markov-switching model produces simulations that can be classified as recessions at a rate consistent with the historical data, whereas the linear model does not.

3.6 Do Expansions Die of Old Age?

The finding in Section 3.3 that a decrease in the unemployment gap in expansions increases the probability of transitioning into recession raises important questions related to the economic literature. Specifically, there is a vast literature on business-cycle duration dependence—that is, whether expansions or recessions "die of old age." The finding above is closely related to that question. In the Markov-switching model, the unemployment rate tends to decline over time in expansions; as the unemployment rate declines, the probability of transitioning from expansion to recession increases. Thus, although the probability of transitioning from expansion to recession is not directly a function of the duration of the expansion, it is indirectly related to the duration of expansion because the unemployment rate tends to decline as the expansion continues.

Diebold and Rudebusch (1990) used nonparametric procedures to test for duration dependence in expansions and recessions, finding some evidence of duration dependence in prewar expansions but little evidence for postwar expansions and recessions. Sichel (1991) conducted a similar exercise using parametric tests, finding statistically significant evidence of positive duration dependence for prewar expansions and postwar recessions. By contrast, Zuehlke (2003) estimated a generalized Weibull hazard model and updated the sample used by Sichel (1991), finding evidence of duration dependence in both postwar recessions and postwar expansions.

Using a Markov-switching framework similar to mine but estimated on GNP growth, Durland and McCurdy (1994) found evidence of duration dependence for recessions but not for expansions. Similarly, Kim and Nelson (1998) investigated duration dependence using a dynamic factor model with regime switching. Using Bayesian estimation methods, they found strong evidence of duration dependence in postwar recessions, but their findings of duration dependence in postwar expansions were sensitive to different priors imposed during estimation.

A principal concern in the estimation of duration dependence in postwar data is the limited sample of expansions and recessions. At the time of Diebold and Rudebusch's (1990) analysis, only nine expansions and recessions had been observed since the end of World War II, severely limiting the power of statistical tests estimated on the data. Moreover, the indirect effect I estimate in the model, with transition probabilities affected by the unemployment gap rather than the duration of expansion or recession, may create distortions when testing for duration dependence.

To investigate the effect of these two factors—limited estimation samples and the indirect link from unemployment gap to duration—on the power of statistical tests to detect duration dependence, I produce simulations from the Markov-switching model and then use one of the nonparametric tests from Diebold and Rudebusch (1990) to see whether this test can identify the indirect effect present in the Markov-switching model. Specifically, I produce 10,000 simulations of the unemployment gap from the model using the procedure discussed in Section 3.5 and then use the Sahm (2019) indicator to classify periods of expansion and

Table 3. Test for Duration Dependence on Simulations

	Expansion	Recession
Mean p Value	0.697	0.864
	(0.241)	(0.182)
Proportion of Simulations With $p < 0.05$	0.003	0.002

Note: The table reports results from a Shapiro-Wilk (1972) test, with null hypothesis that observations are drawn from exponential distribution. Estimates in the table are based on 10,000 samples with nine observations in each sample. Standard deviations are in parentheses.

recession.[16] For each of these classifications, I have a length of duration of the expansion or recession. I then take the first nine expansions and recessions from each of the 10,000 simulations, which gives me 10,000 samples with nine observations of expansions and nine observations of recessions. I then run the Shapiro-Wilk (1972) test on each of these samples.[17] The null hypothesis of this test is that the observations are drawn from the exponential distribution. If the observations are drawn from the exponential distribution, then they will not be duration dependent.[18]

Table 3 displays the results from this procedure. The mean p value for the null hypothesis that the expansions are not duration dependent is 0.697. Moreover, the fraction of simulations in which the null of no duration dependence is rejected at the 5 percent confidence level is just 0.003! Similarly, the mean p value for the null hypothesis that the recessions are not duration dependent is 0.864, and the fraction of simulations in which the null of no duration dependence is rejected at the 5 percent confidence level is just 0.002. This indicates that tests of this nature have almost no power to detect duration dependence when the sample of observations is small and the effect of duration dependence is indirect.

[16] I discard the first 80 quarters of simulations, as well as the first recession and expansion after 80 quarters in each simulation, so that initial conditions do not affect the results.

[17] This test statistic is the W value from Equation 5 and the last column of Table 3 in Diebold and Rudebusch (1990).

[18] See Diebold and Rudebusch (1990) for a thorough discussion of the connection between duration dependence and hazard functions.

4 Forecast Evaluation

In this section, I evaluate the forecasting ability of the benchmark Markov-switching model with TVTPs and compare it with the Markov-switching model with CTPs and the linear version of the model. I evaluate the mean forecasts of the unemployment rate for each model in addition to forecasts of the probability of being in recession.[19] I begin by comparing forecast errors of the models at different horizons over time. I then formally test whether one model outperforms the others using the *GW test* of unconditional predictive ability. I also test whether the relative forecasting performance of the models changes over the business cycle using the *GR test*, which evaluates the evolution of the models' relative forecast performance over time. I find that the benchmark TVTP Markov-switching model produces smaller out-of-sample RMSEs than either alternative model beyond the 10-quarter forecast horizon with varying degrees of statistical significance.

4.1 Forecast Comparison of Linear AR(2) Model, TVTP Markov-Switching Model, and CTP Markov-Switching Model

I compare the forecasting performance of three competing models over time: the linear AR(2) model of the unemployment gap, the benchmark Markov-switching model of the unemployment gap with TVTPs, and the Markov-switching model of the unemployment gap with CTPs.[20] For consistency with the GW testing framework, I estimate the models over rolling samples of 120 quarters, with the first sample including data from Q1 1959 through Q4 1988 and the last sample including data from Q1 1985 through Q4 2014, giving me a sample of 104 forecasts.[21] To construct the forecasts for each estimated model, I produce 10,000 simulations

[19] To identify periods that may be classified as recessions, I use the Sahm indicator described in Section 3.5.

[20] As pointed out by a reviewer, when forecasting a stationary time series, at some forecast horizon no model should be able to outperform the sample average of the series being forecast. I explored this using the models and forecasting procedure and found that this occurs beyond the five-year forecast horizon I consider below.

[21] This end date was selected so that the five-year forecast errors did not include data from the COVID-19 recession. The relative forecast performance of each model at longer-term horizons is unchanged if these data points are included, though the RMSE for each model is greater at longer forecast horizons.

Figure 8. Forecast Errors of Mean Unemployment Rate

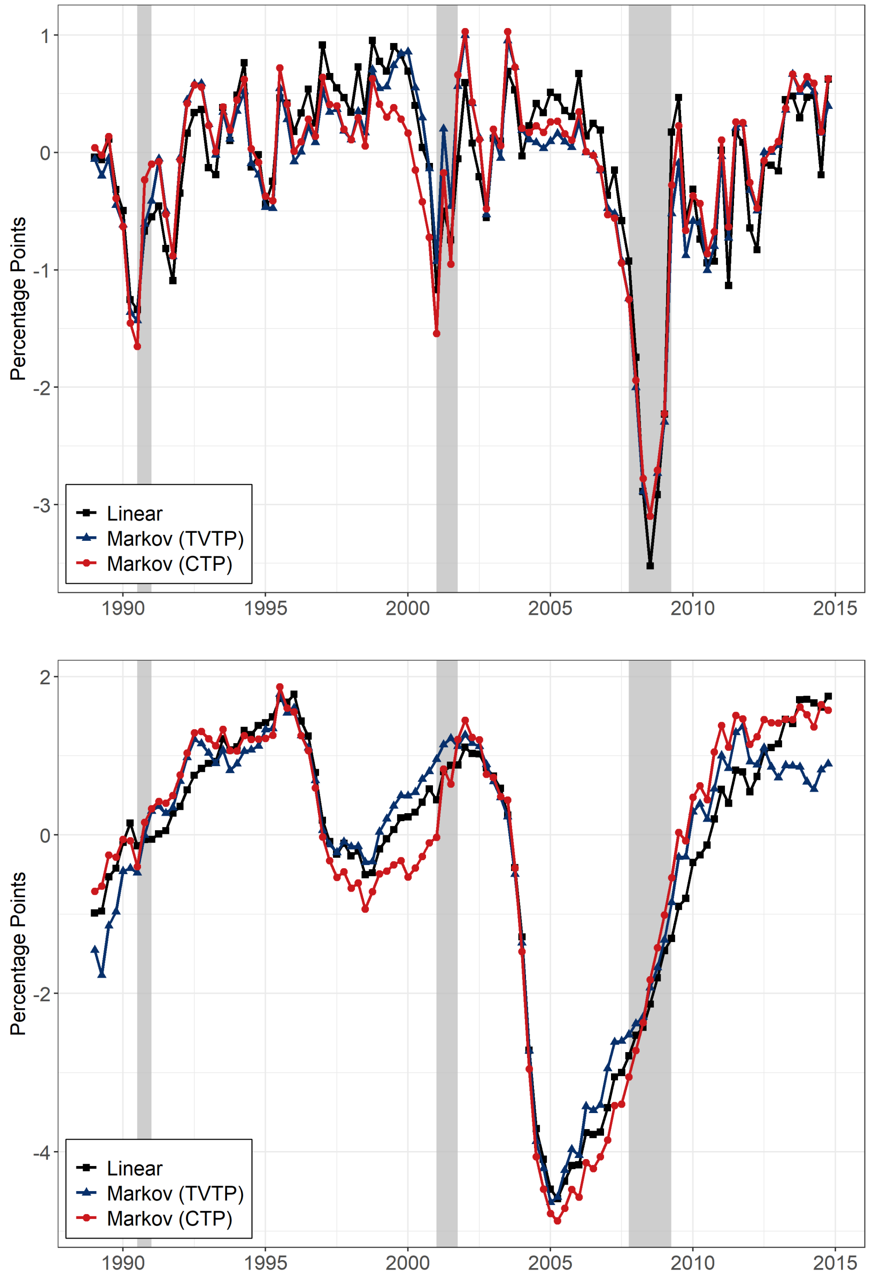

Note: The top panel plots forecast errors at a 4-quarter horizon for forecasts made at time t; the bottom panel plots forecast errors at a 20-quarter horizon for forecasts made at time t. Gray bars denote periods identified as recessions by the National Bureau of Economic Research. CTP = constant transition probabilities; TVTP = time-varying transition probabilities.

Figure 9. Root Mean Square Errors Across Forecast Horizons

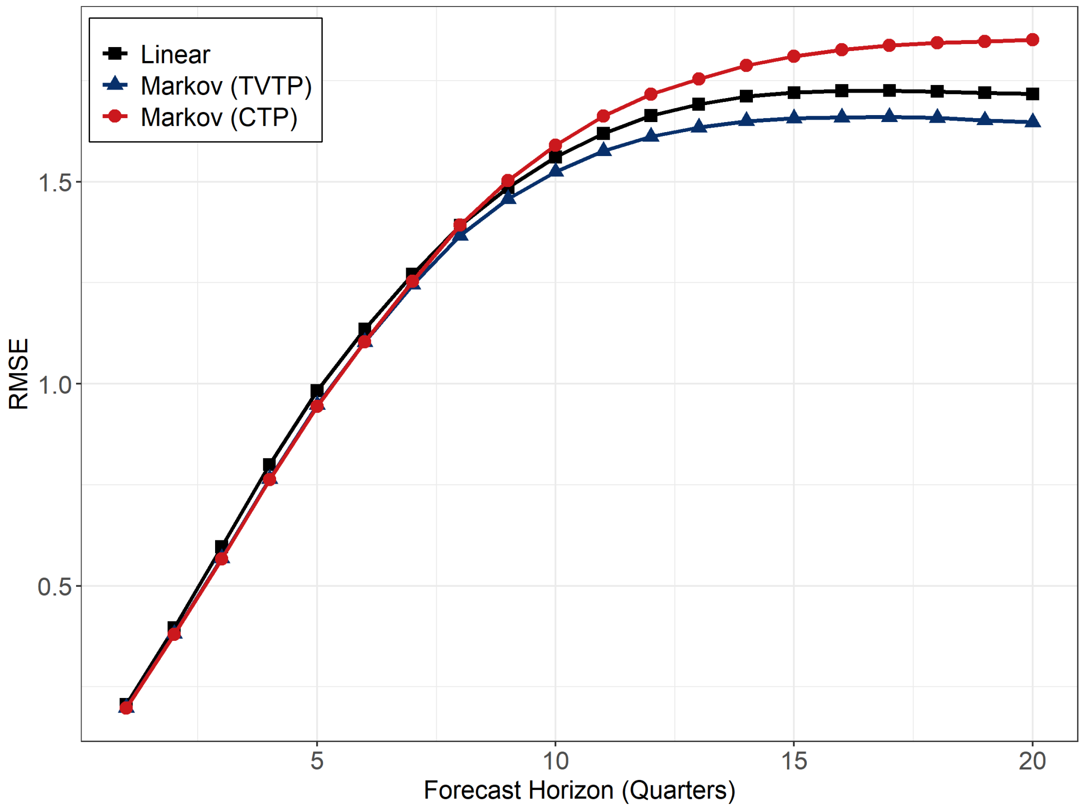

Note: CTP = constant transition probabilities; RMSE = root mean square error; TVTP = time-varying transition probabilities.

with the shocks drawn from the sample of estimated errors. For the Markov-switching models, the procedure for producing the simulations is identical to the procedure described in Section 3.5. This produces 10,000 paths of the unemployment rate for each model, with the mean forecast simply taking the average across the paths. For the forecast of the probability of being in recession, I calculate the fraction of simulations that activate the Sahm indicator described in Section 3.5.[22] I construct forecasts over a 20-quarter horizon, which allows me to compare the forecasts from one quarter to five years into the future.

The top panel of Figure 8 shows the forecast errors of the mean unemployment rate for each model at the four-quarter horizon, with the date on the x-axis indicating the date at

[22]This can be thought of as an out-of-sample version of the exercise presented in Section 3.5.

which the forecast was made. On casual inspection, it is somewhat unclear which model produces smaller absolute forecast errors at the 4-quarter horizon on average. To compare the forecasts over a longer time horizon, the bottom panel of Figure 8 shows the forecast error for each model at the 20-quarter horizon. At this forecast horizon, it appears that the linear model produced somewhat superior forecasts for years prior to 2000, whereas the TVTP Markov-switching model produced slightly superior forecasts thereafter.

To compare the average forecasting performance of the models across forecast horizons, Figure 9 shows the RMSE of the mean unemployment rate for each model from 1 quarter to 20 quarters into the future. The models have similar RMSEs over the first several quarters of the forecasts, with the Markov-switching models producing slightly smaller RMSEs on average. As the forecast horizon increases, however, the relative forecast performance of the CTP Markov-switching model gradually deteriorates, and the TVTP Markov-switching model outperforms both the CTP Markov-switching model and the linear model. Of course, from this analysis alone we cannot conclude that one forecasting model is better than another. For that we require a formal test of whether the differences in squared forecast errors are statistically significant.

4.2 Giacomini-White Test of Unconditional Predictive Ability

I use the GW test of unconditional predictive ability to test whether the differences in squared forecast errors are statistically significant. The GW test of unconditional predictive ability tests the null hypothesis that the average difference in squared forecast errors between two models is equal to zero. This null hypothesis states that we cannot predict which forecasting method will be more accurate at a given forecast horizon. This test involves estimating each model over a rolling sample and producing forecasts for each model estimated over the rolling sample. The use of a rolling estimation window leads to estimators that converge asymptotically to standard probability distributions, which allows for easy comparisons between models.[23]

I estimate the models and produce forecasts using the same rolling sample windows and forecast horizons as in Section 4.1. The top panel of Figure 10 displays the one-sided GW test statistic for the mean unemployment rate, comparing the forecasting ability of the TVTP Markov-switching model with that of the linear model across forecast horizons.[24] In this comparison, a negative value of the test statistic indicates that the squared forecast errors are greater for the TVTP Markov-switching model, whereas a positive value of the test statistic indicates that the squared forecast errors are greater for the linear model. Consistent with the analysis of the RMSEs for each model, the TVTP Markov-switching model produces squared forecast errors smaller than those produced by the linear model at forecast horizons up to 20 quarters. Additionally, the difference in squared forecast errors is statistically significant at the 5 percent level at the 1-quarter horizon and is statistically significant at the 10 percent level at the 2- and 3-quarter horizons and most forecast horizons greater than 13 quarters. From this we can infer that, with some degree of statistical significance, the forecasts produced by the TVTP Markov-switching model outperform those produced by the linear model at forecast horizons less than a year and at forecast horizons greater than 13 quarters.

To show how the use of TVTPs affects the forecasting ability of the Markov-switching model, the bottom panel of Figure 10 displays the one-sided GW test statistic comparing the forecasting ability of the TVTP Markov-switching model with that of the CTP Markov-switching model. The CTP model produces squared forecast errors that are smaller at forecast horizons of less than 6 quarters and greater thereafter. Moreover, the difference in squared forecast errors is statistically significant at the 10 percent critical value 10 to 13 quarters into the future and is statistically significant at the 5 percent critical value at forecast horizons greater than 13 quarters. This suggests that at longer forecast horizons, the use of TVTPs significantly improves the Markov-switching model's forecast performance.

I now compare the ability of each model to forecast the probability of being in recession

[23] Tests that compare nested models as in Diebold and Mariano (1995) and West (1996) do not converge in distribution to standard probability distributions, which requires simulation to calculate critical values.

[24] I use squared error loss for the loss function in the comparison of the forecast errors.

Figure 10. GW Statistics for Mean Unemployment-Rate Forecasts

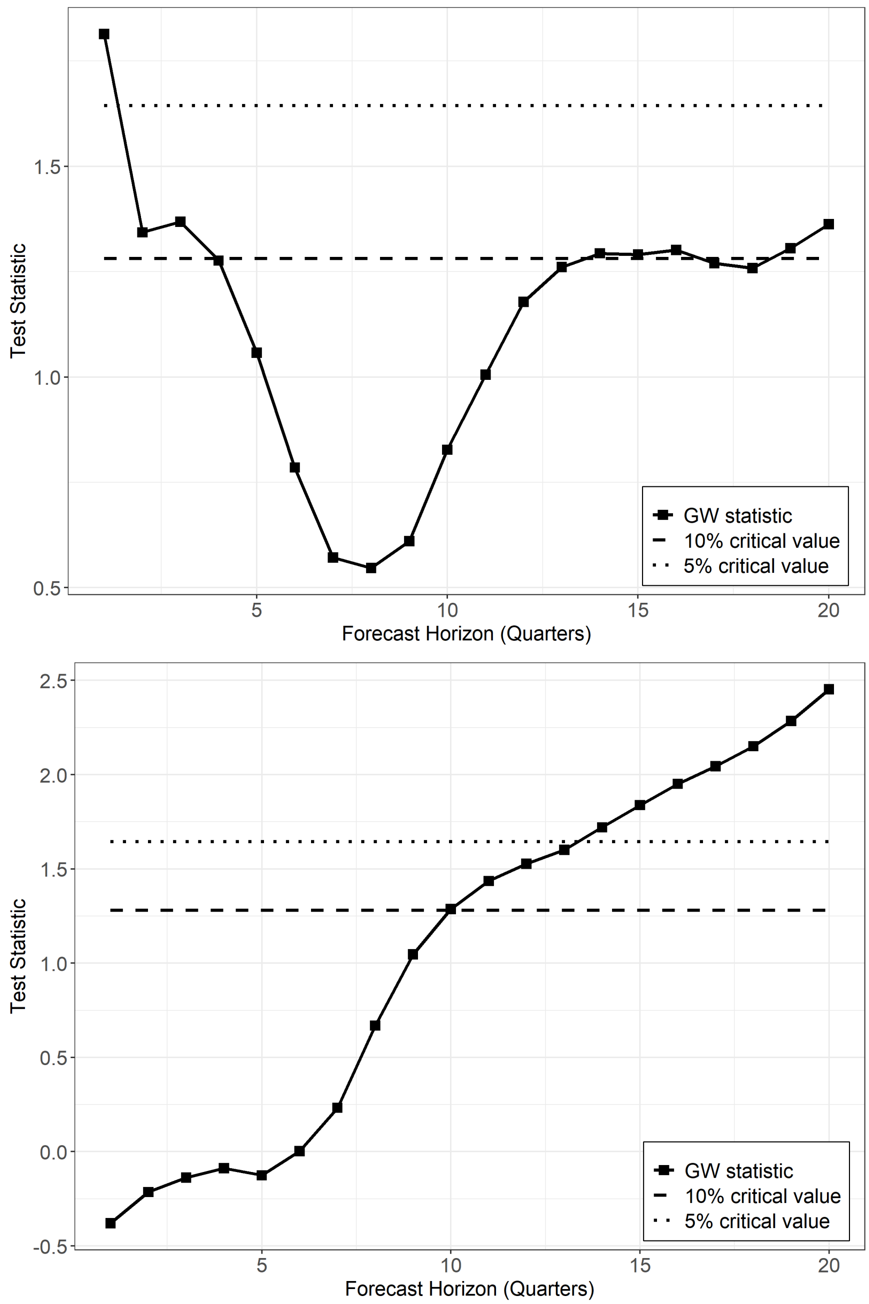

Note: The top panel plots GW test statistics comparing forecasts from the linear AR(2) model with those from the TVTP Markov-switching model. The bottom panel plots GW test statistics comparing forecasts from the CTP Markov-switching model with those from the TVTP Markov-switching model. In both panels, a positive test statistic indicates that the TVTP Markov-switching model outperforms the alternative model. CTP = constant transition probability; GW = Giacomini-White; TVTP = time-varying transition probability.

Figure 11. GW Statistics for Recession-Probability Forecasts

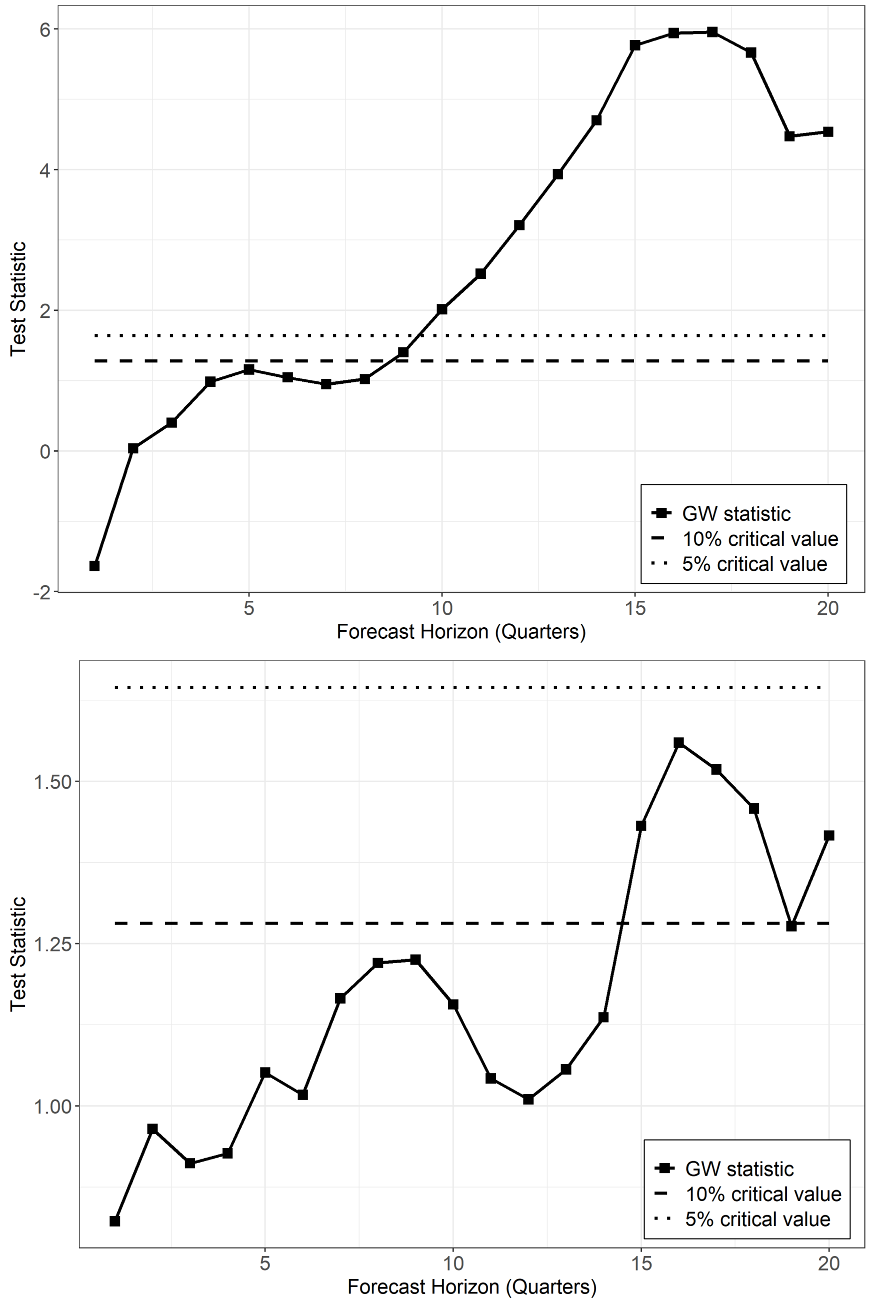

Note: The top panel plots GW test statistics comparing forecasts from the linear AR(2) model with those from the TVTP Markov-switching model. The bottom panel plots GW test statistics comparing forecasts from the CTP and TVTP Markov-switching models. In both panels, a positive test statistic indicates that the TVTP Markov-switching model outperforms the alternative model. CTP = constant transition probability; GW = Giacomini-White; TVTP = time-varying transition probability.

as identified by the Sahm indicator. The top panel of Figure 11 displays the one-sided GW test statistic for the probability of activating the Sahm indicator, comparing the forecasting ability of the TVTP Markov-switching model with that of the linear model across forecast horizons.[25] The linear model appears to produce smaller squared forecast errors at the 1-quarter horizon but larger squared errors thereafter. Additionally, the difference in squared errors is statistically significant at the 10 percent level at the 9-quarter horizon and at the 5 percent level thereafter. This suggests that at longer forecast horizons, the Markov-switching model produces superior forecasts of the probability of being in recession relative to the linear model with a high degree of statistical significance.

To illustrate how the use of TVTPs affects the ability of the Markov-switching model to forecast the probability of recession, the bottom panel of Figure 11 displays the one-sided GW test statistic comparing the forecasting ability of the TVTP Markov-switching model with that of the CTP Markov-switching model. The CTP model produces squared forecast errors larger than those produced by the TVTP model at each forecast horizon. However, these differences in squared forecast errors are statistically significant at the 10 percent critical value only at forecast horizons of 14 quarters or more. Nonetheless, this provides some evidence that at longer forecast horizons, the TVTP Markov-switching model produces better forecasts of the probability of being in recession than the CTP Markov-switching model.

4.3 Giacomini-Rossi Fluctuation Test

Given that the Markov model predicts differing dynamics over the business cycle, I am also interested in whether it produces better forecasts in expansions or recessions. The GR fluctuation test allows me to answer this question. This test can be thought of as a local version of the global GW test. In the GW test, the estimation is performed using the entire sample of forecast errors. In the GR test, the estimation is performed on a rolling subsample of

[25] Although the linear model is not designed to forecast recessions, it is still useful as a benchmark for producing simulations that may be characterized as recessionary.

Figure 12. GR Statistics for Mean Unemployment-Rate Forecasts

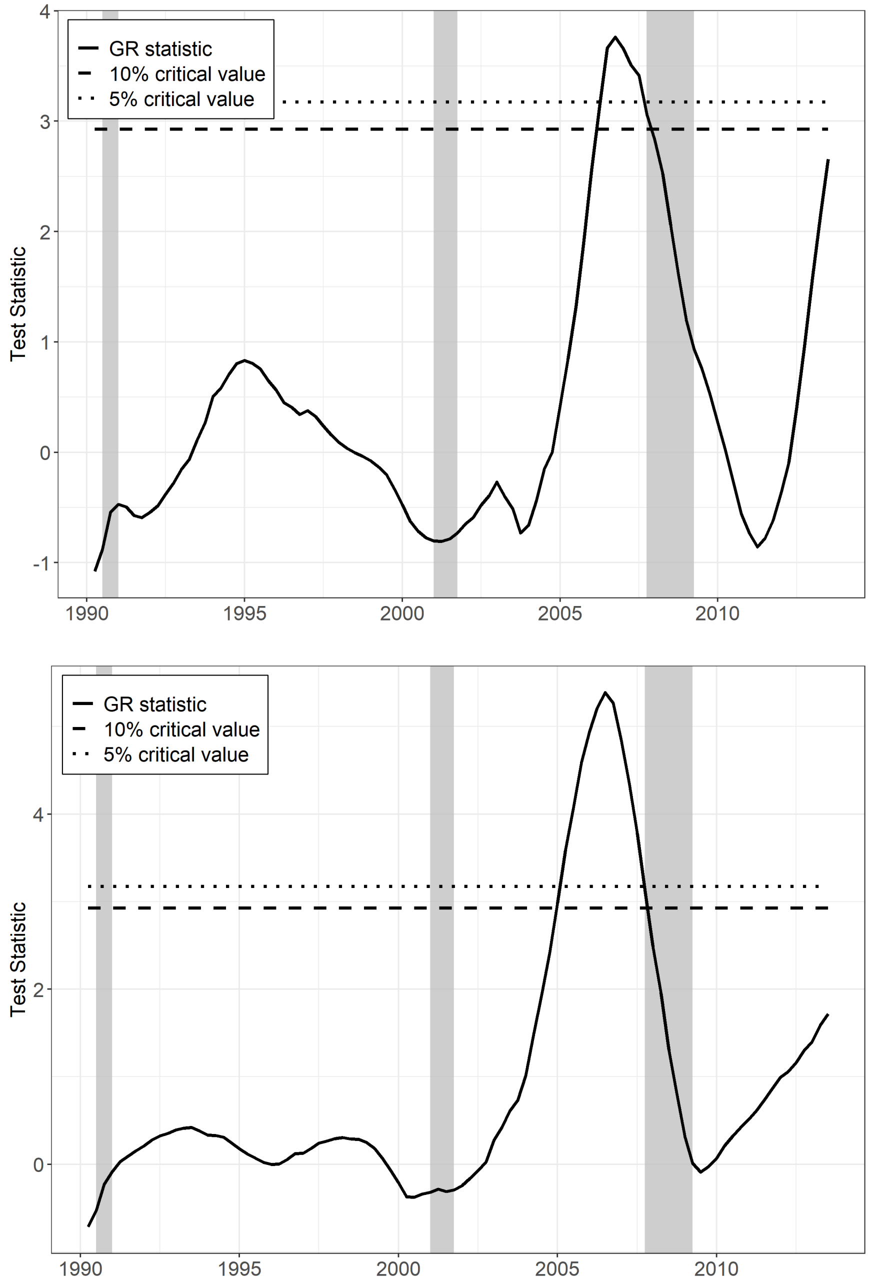

Note: The top panel plots GR test statistics comparing forecasts from the linear AR(2) model and the TVTP Markov-switching model; the bottom panel plots GR test statistics comparing forecasts from the CTP and TVTP Markov-switching models. In both panels, a positive test statistic indicates that the TVTP Markov-switching model outperforms the alternative model. Gray bars denote periods identified as recessions by the National Bureau of Economic Research. CTP = constant transition probability; GR = Giacomini-Rossi; TVTP = time-varying transition probability.

the forecast errors, which allows for inferences about how the relative forecast performance of each model changes over time. Because the forecasts from the TVTP Markov-switching model outperform those produced by both the linear model and the CTP Markov-switching model at longer forecast horizons, I evaluate the forecasts at the five-year horizon. Additionally, I use 10 quarters of data for the rolling subsamples of the forecast errors.[26]

The top panel of Figure 12 displays the one-sided GR test statistic for the mean unemployment rate, comparing the local forecasting ability of the TVTP Markov-switching model with the forecasting ability of the linear model over time. The values in the GR test are interpreted similarly to those in the GW test, with negative values indicating that the squared forecast errors are greater for the TVTP Markov-switching model and positive values indicating that the squared forecast errors are greater for the linear model. Although the values have a similar interpretation, the GR test statistic does not converge in distribution to a standard probability distribution, so the critical values do not match those of the GW test. Consistent with the analysis of the forecast errors over time, the relative forecasting performance of the TVTP Markov-switching model has fluctuated over time, especially prior to 2005. However, the forecasts from that model dramatically outperformed those from the linear model from 2005 to 2010, with the test statistic surpassing the 5 percent critical value just before the 2007–2009 recession. The relative forecasting performance of the Markov-switching model has also increased dramatically since 2013, but not to a statistically significant extent. These results reflect that the TVTP Markov-switching model produced a mean forecast for the unemployment rate greater than that produced by the linear model from 2005 to 2010 and less than that produced by the linear model after 2013.

To illustrate how the use of TVTPs affects the ability of the Markov-switching model to forecast the mean unemployment rate over time, the bottom panel of Figure 12 displays the one-sided GR test statistic comparing the local forecasting ability of the TVTP Markov-switching model with the forecasting ability of the CTP Markov-switching model over time. As in the previous comparison, the TVTP model outperformed the CTP model at a statisti-

[26] This corresponds to the 0.1 value of μ in the critical values tabulated in Giacomini-Rossi (2010).

Figure 13. GR Statistics for Recession-Probability Forecasts

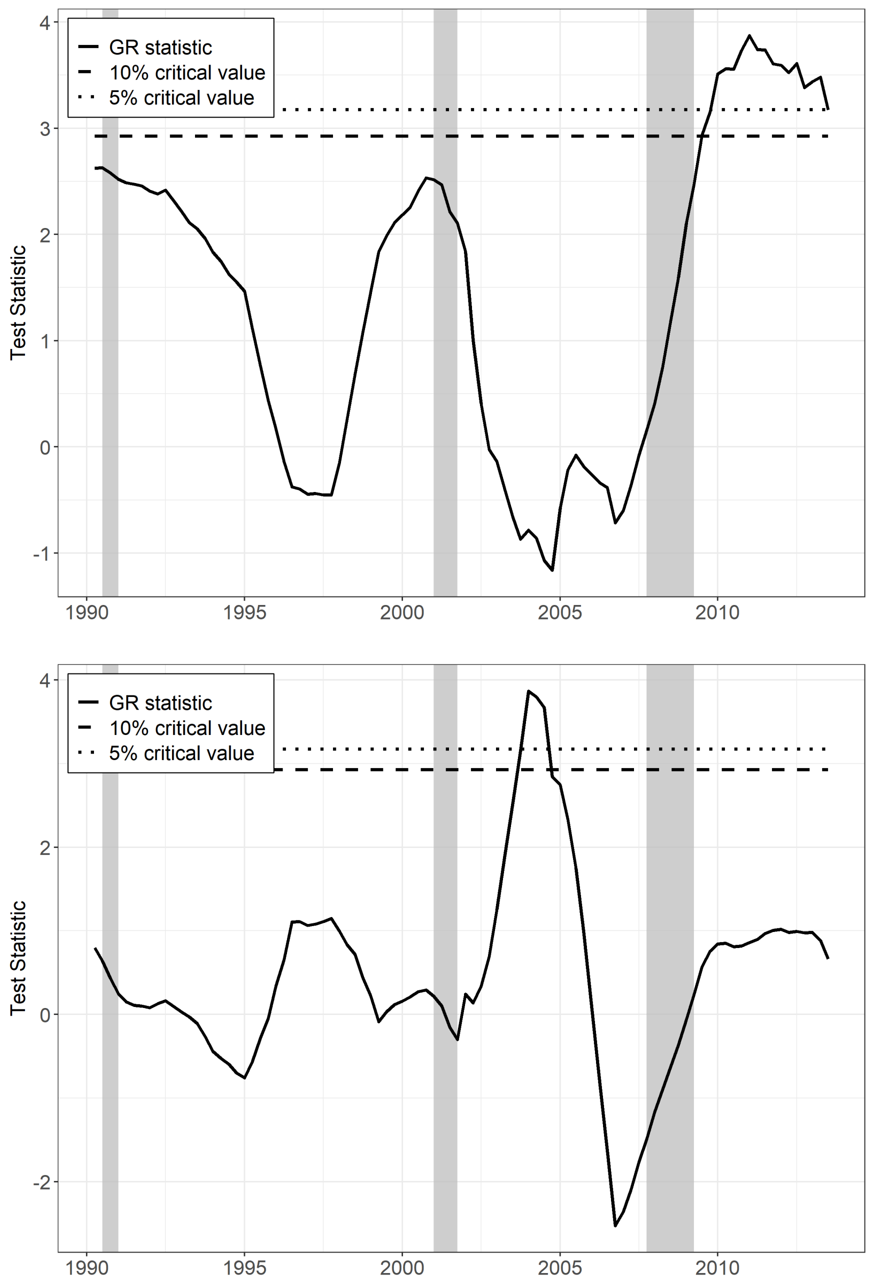

Note: The top panel plots GR test statistics comparing forecasts from the linear AR(2) model with those from the TVTP Markov-switching model. The bottom panel plots GR test statistics comparing forecasts from the CTP and TVTP Markov-switching models. In both panels, a positive test statistic indicates that the TVTP Markov-switching model outperforms the alternative model. Gray bars denote periods identified as recessions by the National Bureau of Economic Research. CTP = constant transition probability; GR = Giacomini-Rossi; TVTP = time-varying transition probability.

cally significant level before the 2007–2009 recession and at a somewhat lesser rate thereafter. This suggests that the use of TVTPs has increased the relative forecasting performance of the Markov-switching model over the past two decades.

I now compare how the ability of each model to forecast the probability of being in recession has changed over time using the GR statistic. The top panel of Figure 13 displays the one-sided GR test statistic for the probability of activating the Sahm indicator, comparing the local forecasting ability of the TVTP Markov-switching model with the forecasting ability of the linear model over time. The local relative forecasting ability of the TVTP Markov-switching model has fluctuated considerably over time but is statistically significant at the 5 percent critical value following the 2007–2009 recession. The TVTP Markov-switching model produced very few simulations characterized as recessions during this period (when the unemployment rate was high), whereas the linear model continued to produce many simulations characterized as recessions. Given that no recession occurred within five years after the 2007–2009 recession, the TVTP Markov-switching model was more accurate at forecasting the probability of being in recession during this period.

The bottom panel of Figure 13 displays the one-sided GR test statistic comparing the local forecasting ability of the TVTP Markov-switching model with the forecasting ability of the CTP Markov-switching model over time. As in the previous comparison, the local relative forecasting ability of the TVTP Markov-switching model has fluctuated over time and is statistically significant at the 5 percent level just before 2005. This is a result of the low unemployment gap, which leads the TVTP Markov-switching model to produce more simulations that enter recession within the next five years than the CTP Markov-switching model. Because the TVTP Markov-switching model produced more recession simulations just before the 2007–2009 recession, the use of TVTPs made that model more accurate than the model with CTPs during this period.

5 Robustness

In this section I consider two robustness checks for the results presented in Section 3. First, I consider the sensitivity of the results to the use of TVTPs. Second, I examine how the dynamics of the model change when the unemployment rate is used in place of the unemployment gap. To summarize, neither of these modeling assumptions has a significant effect on the results presented in Section 3.

5.1 Time-Varying or Constant Transition Probabilities?

It is possible that I am overfitting the model by using TVTPs, which may be affecting the underlying dynamics of the model. Some of the coefficients on the lagged unemployment gap in the transition probability equations are imprecisely estimated, suggesting that the use of TVTPs may not be warranted. To investigate how TVTPs affect the results in the preceding section, I reestimate the model using CTPs, dropping the lagged unemployment gap as an explanatory variable in the transition probability equations. I make no other changes to the model and use the same sample for estimation.

Panel B of Table 1 displays the estimated parameters of the model, along with the associated standard errors. To compare the dynamics of the model estimated using CTPs, the bottom panel of Figure 15 displays the deviation of the unemployment gap from the model's steady state after transitioning to each state for at least one quarter. Transition to the recession state is characterized by an increase in the unemployment gap of slightly more than 0.8 percentage points over eight quarters, almost identical to the increase in the TVTP model. Similarly, transition to the expansion state is characterized by a decrease in the unemployment gap of almost 0.4 percentage points, similar to the decrease in the TVTP model. One noticeable difference between the TVTP and CTP models is the steady-state unemployment gap, which is roughly 0.5 percentage points in the TVTP model and 0.1 percentage point in the CTP model.

Because transition probabilities are constant, they do not vary with the unemployment

Figure 14. Probabilities of Recession Using Constant Transition Probabilities

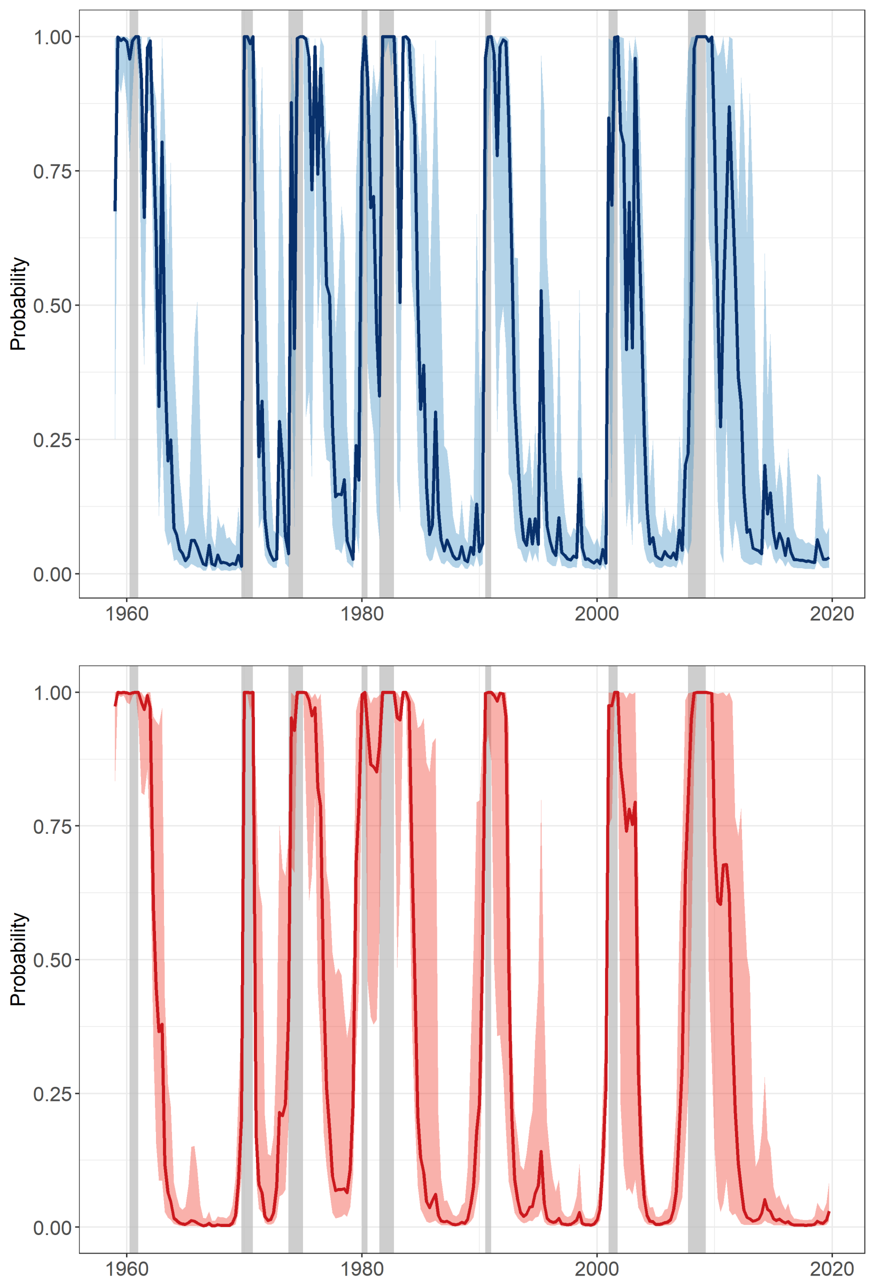

Note: The top panel plots the filtered probability of being in a recession state at time t; the bottom panel plots the smoothed probability of being in a recession state at time t. Shaded regions are 95 percent confidence intervals from parametric bootstrap. Gray bars denote periods identified as recessions by the National Bureau of Economic Research.

gap and are instead fixed across the business cycle. The estimated probability of staying in the recession state is 0.902, with a 95 percent confidence interval of 0.786 to 0.958. Meanwhile, the estimated probability of staying in the expansion state is 0.950, with a 95 percent confidence interval of 0.893 to 0.978.

Figure 14 displays the weight that the CTP model attaches to the recession state over time. These estimates are nearly identical to those presented in Section 3 and also correspond closely to NBER recessions.

5.2 Unemployment Gap or Unemployment Rate?

One potential issue with the preceding analysis is the use of the unemployment gap rather than the unemployment rate. That is, it could be the case that the estimated model dynamics are merely an artifact of de-trending the unemployment rate by CBO's estimate of the noncyclical rate of unemployment rate. Because the noncyclical rate of unemployment is unobserved, it must be estimated using various econometric methods. These methods can yield differing estimates of the noncyclical rate, complicating analyses using the unemployment gap. To examine this possibility, I reestimate the model, replacing the unemployment gap with the unemployment rate in each equation. The model is estimated over the same sample used in the previous section.

Panel C of Table 1 displays the estimated parameters of the model, along with the associated standard errors. To compare the dynamics of the model estimated with the unemployment rate, the top panel of Figure 15 displays the deviation of the unemployment rate from the model's steady state after transitioning to each state for at least one quarter. This figure corresponds to Figure 4 in the previous section and can be interpreted similarly. Transition to the recession state is characterized by an increase in the unemployment rate of almost 1 percentage point over eight quarters, compared with an increase of slightly more than 0.8 percentage points in the model estimated with the unemployment gap. By contrast, transition to the expansion state is characterized by a decrease in the unemployment rate of roughly 0.4 percentage points over eight quarters, nearly identical to the decline for the

Figure 15. Dynamics of Unemployment After Transition Using the Unemployment Rate and Constant Transition Probabilities

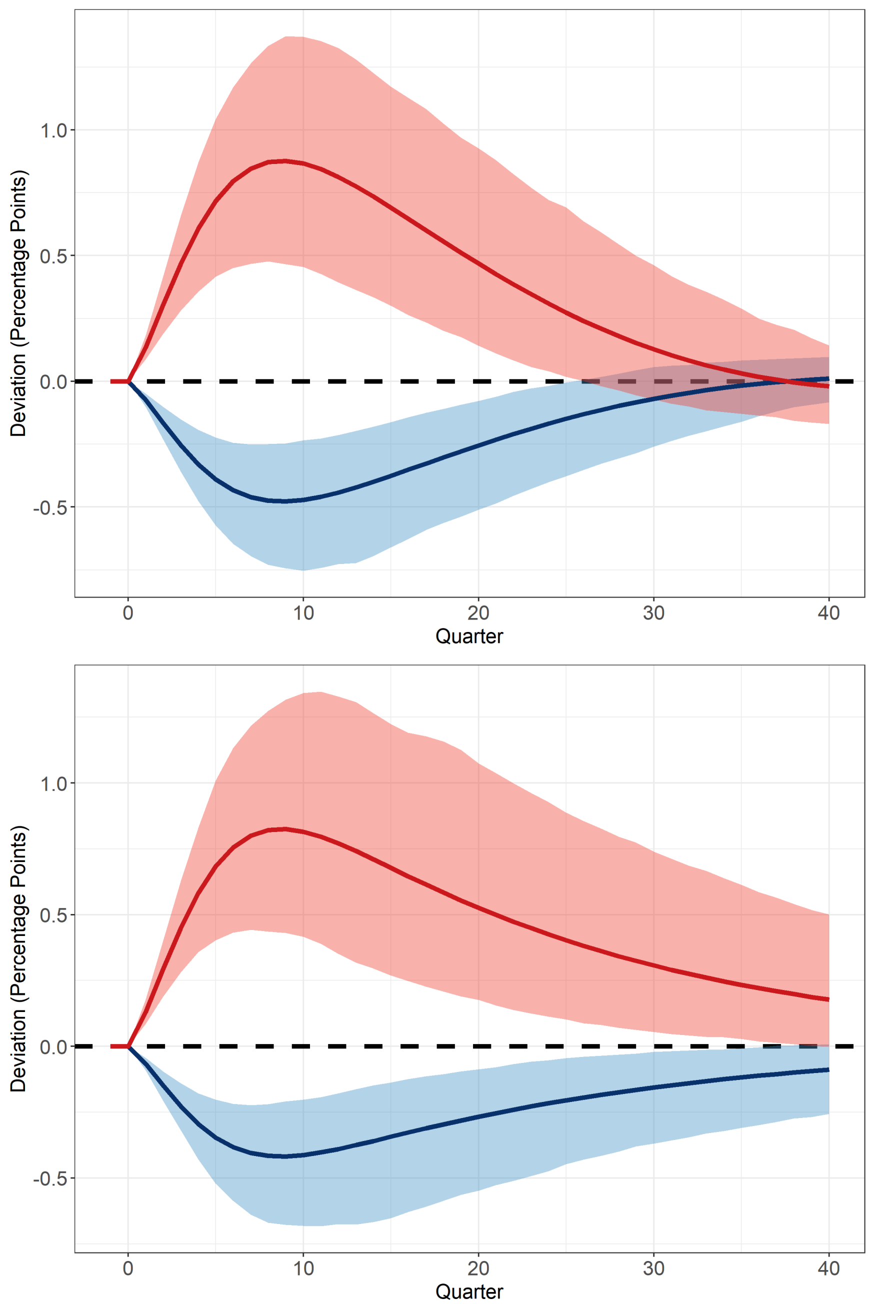

Note: The top panel plots the deviation of the unemployment rate from the model's steady state after transition to each state for at least one quarter. The bottom panel plots the deviation of the unemployment gap from the model's steady state after transition to each state for at least one quarter. Blue denotes expansion states; red denotes recession states. Solid lines are averages across 1,000 simulations with no shocks using estimated parameters. Shaded regions are 95 percent confidence intervals from parametric bootstrap.

Figure 16. Quarterly Transition Probabilities Using the Unemployment Rate

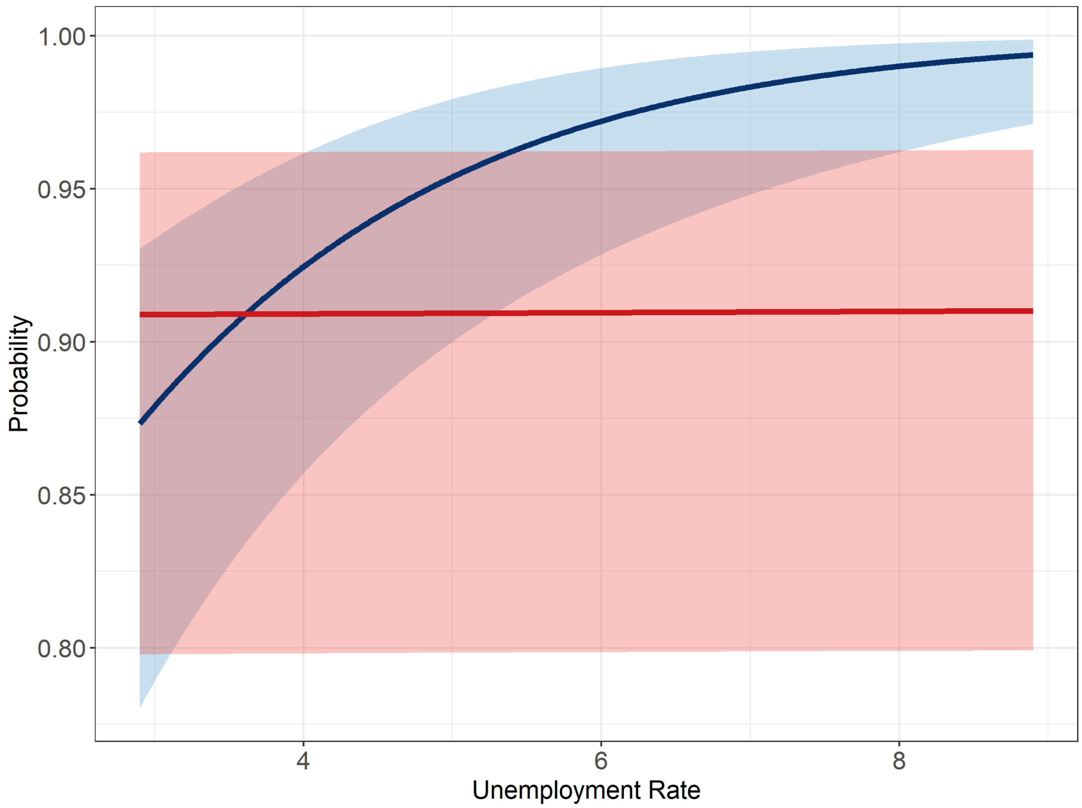

Note: The graph plots the probability of transitioning from state i at time t to state i in time $t+1$ conditional on the unemployment rate. Blue denotes expansion states; red denotes recession states. Solid lines are transition probabilities using point estimates of parameters. Shaded regions are 95 percent confidence intervals from parametric bootstrap.

model estimated with the unemployment gap.

To illustrate the effect on transition probabilities, Figure 16 displays the estimated transition probabilities as a function of the unemployment rate. The probability of staying in the expansion state declines as the unemployment rate declines, similar to the result presented in the previous section. Moreover, this effect is statistically significant. By contrast, the probability of staying in the recession state is roughly constant as the unemployment rate changes, though this parameter is estimated much less precisely than in the expansion state.

Figure 17 displays the weight that the model estimated with the unemployment rate

Figure 17. Probabilities of Recession Using the Unemployment Rate

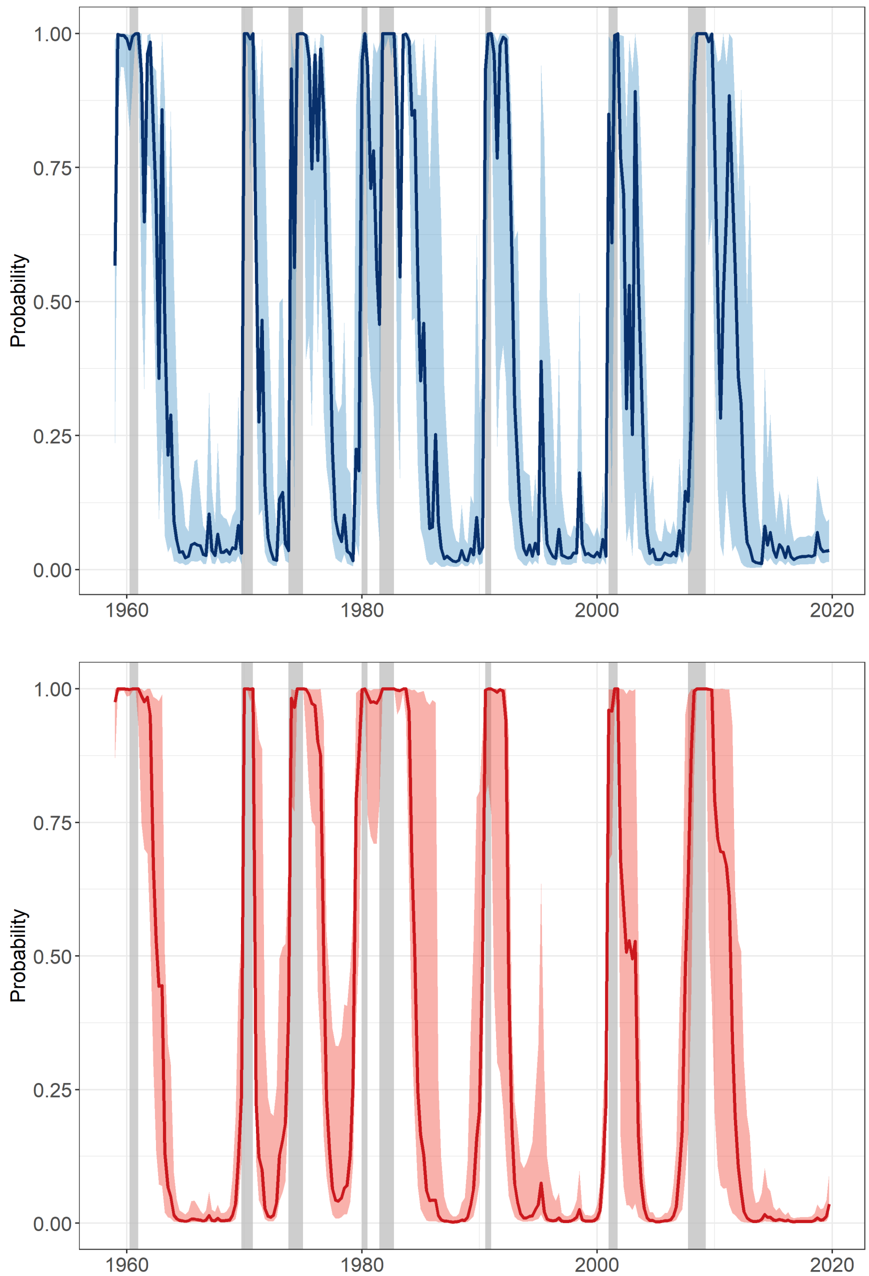

Note: The top panel plots the filtered probability of being in a recession state at time t; the bottom panel plots the smoothed probability of being in a recession state at time t. Shaded regions are 95 percent confidence intervals from parametric bootstrap. Gray bars denote periods identified as recessions by the National Bureau of Economic Research.

attaches to the recession state over time. Again, these estimates are nearly identical to those presented in Section 3, and also correspond closely to NBER recessions.

6 Pandemic Recession and Simulations

In this section, I document the adjustments made to account for the significant changes in the unemployment-rate data since the start of the COVID-19 pandemic. Specifically, I split the unemployment data into two components: permanently separated unemployment and temporarily separated unemployment. I then model these components individually, using a Markov-switching model for the permanently separated unemployment rate (PSUR) and a standard linear time-series model for the temporarily separated unemployment rate (TSUR). I then discuss the procedure for creating simulations of the unemployment rate using the models for the PSUR and TSUR, as well as the unemployment-gap model presented in Section 3.

6.1 Permanently Versus Temporarily Separated Unemployment

The recession and recovery associated with the COVID-19 pandemic present unique challenges for modeling time series of economic data. In April 2020, the unemployment rate increased by 10.3 percentage points, nearly 50 times the prepandemic standard deviation of 0.21 for the monthly change of the unemployment rate. Moreover, most of the increase in the unemployment rate was driven by workers classified as temporarily separated unemployed. Workers are classified as temporarily separated unemployed if they have experienced a temporary separation from their employer and expect to be recalled to work by that employer. By contrast, workers who experience a permanent separation from their employer are classified as permanently separated unemployed. (For clarification, I do not mean that a worker who is permanently separated unemployed will remain unemployed in perpetuity. Rather, I mean that the separation from their previous employer is permanent and that they will remain unemployed until they either find employment with a new employer or exit the labor force.)

Figure 18 displays the PSUR, defined as the fraction of the labor force classified as permanently separated unemployed, and the TSUR, defined as the fraction of the labor force classified as temporarily separated unemployed. As is clear from the figure, most of the unprecedented increase in the unemployment rate in April 2020 was driven by an increase in workers classified as temporarily separated unemployed. This stands in stark contrast to previous recessions, in which the bulk of the increase in the unemployment rate was driven by increases in the number of permanently separated unemployed workers. As noted by Gallant et al. (2020), temporarily separated unemployed workers return to employment at much higher rates than permanently separated unemployed workers, which accounts for the rapid decline in the TSUR in Figure 18. Moreover, most of the workers classified as unemployed at the start of the recession were recorded as temporarily separated unemployed. Given these unique circumstances, the Markov-switching model of the unemployment gap presented in the previous sections would likely provide misleading simulations in the current environment, because it characterizes recoveries as gradual declines in the unemployment rate.

Despite the extreme dynamics of the TSUR since the start of the COVID-19 recession, the PSUR has for the most part followed a pattern typical of past recessions.[27] As noted above, increases in the unemployment rate in past recessions were primarily driven by increases in the number of permanently separated unemployed workers. As a result, the PSUR appears to exhibit the same dynamics discussed in Section 2, increasing rapidly in recessions and falling gradually in expansions. It is less clear if the TSUR exhibits this asymmetric dynamic. To formally test whether both the PSUR and the TSUR exhibit the asymmetric dynamic discussed in Section 2, I run the Bai-Ng (2005) test for skewness on the change in each unemployment rate, along with the headline unemployment rate for comparison. This test involves estimating the coefficient of skewness for a given variable, then testing the null hypothesis that there is no skewness in a variable by using a covariance matrix that is robust to serial correlation in the data. In this case, if the change in an unemployment rate is

[27] The behavior of the PSUR during 2021 was somewhat anomalous, declining at a faster rate than in past expansions.

Figure 18. Permanently and Temporarily Separated Unemployment Rates

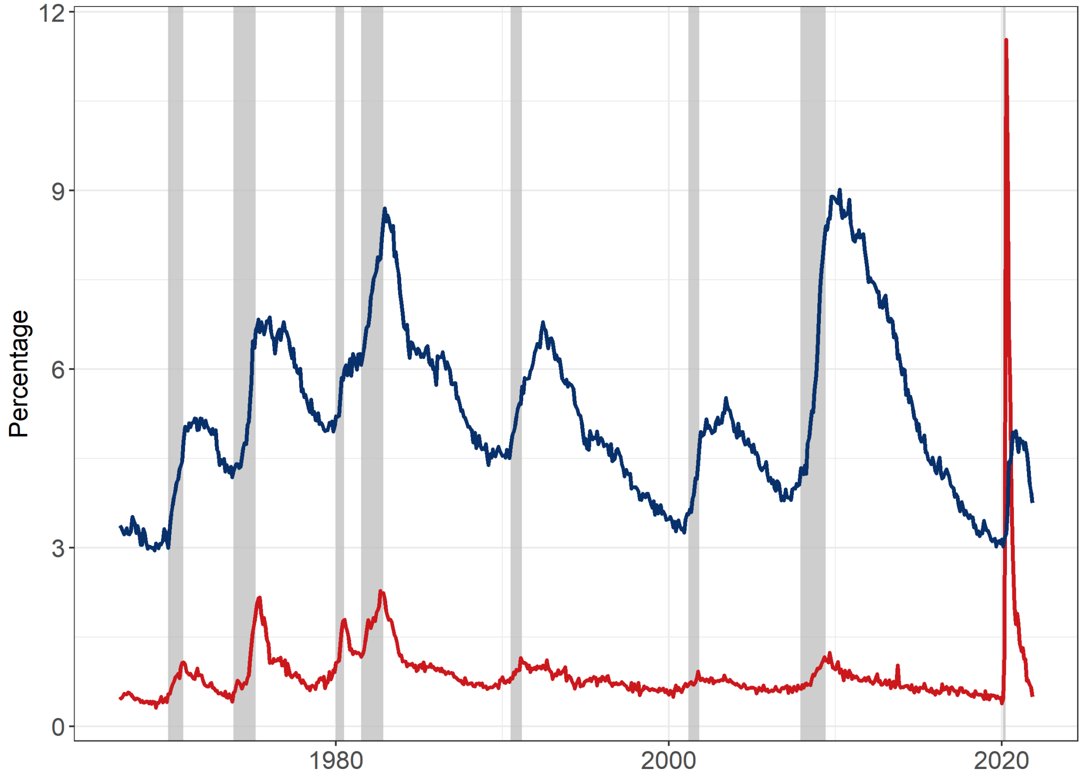

Note: The solid blue line denotes the permanently separated unemployment rate; the solid red line denotes the temporarily separated unemployment rate.

positively skewed, then that unemployment rate increases at a faster rate than it declines.

Table 4 shows the coefficient of skewness for each variable using monthly data over the sample from February 1967 through December 2019.[28] Consistent with the literature reviewed in Section 2.1, the change in the headline unemployment rate is positively skewed with a high degree of statistical significance. The change in the PSUR is also positively skewed at a statistically significant level, matching the observation above that the PSUR appears to increase at a faster rate than it declines. Alternatively, the change in the TSUR has a coefficient of skewness very close to zero, and we cannot reject the null hypothesis of zero skewness with a reasonable level of statistical significance. Because of these findings,

[28]Historical data on the split between TSUR and PSUR are available only back to 1967.

Table 4. Skewness of Change in Unemployment Rates

Δ	Skewness
Unemployment Rate	0.663**
	(.032)
Permanently Separated Unemployment Rate	0.477**
	(.018)
Temporarily Separated Unemployment Rate	0.027
	(.461)

Note: The table reports coefficients of skewness for change in unemployment rates, with p values from one-sided Bai-Ng (2005) tests in parentheses, over the sample from February 1967 to December 2019. $^{**}p < 0.05$.

I estimate the Markov-switching model using the PSUR while estimating a simple linear model for the TSUR.[29] I can then combine the PSUR and TSUR simulations to produce simulations of the headline unemployment rate.[30] Finally, once TSUR simulations return to more normal levels, I can use the Markov-switching model of the quarterly unemployment gap to produce simulations of the headline unemployment rate.

I estimate the Markov-switching model described in Section 3 with two changes. First, I estimate the model using monthly PSUR data over the sample spanning April 1967 through December 2019.[31] Second, I use three lags of the dependent variable as explanatory variables instead of two, keeping only one lag of the dependent variable as an explanatory variable in the equations governing transition probabilities. Although the Akaike information criterion and Bayes/Schwarz information criterion of a linear AR process for the PSUR are minimized at higher-order lags, I use just three lags to keep the estimation tractable.[32] Figure 19 shows the estimated probability of being in the recession state at each point in history.

[29] Note that I am not explicitly modeling labor market flows, but rather the rates that are a result of these flows.

[30] Because the denominator in each unemployment-rate series is the total labor force, I can simply sum the two unemployment rates to create the headline unemployment rate.

[31] Our data set for the split between the PSUR and TSUR begins in January 1967.

[32] Estimation of the Markov-switching model incorporates the assumption that the conditional density depends only on the current state, not on past states. To implement this assumption, the number of states used when estimating the model increases exponentially with the number of lags.

Figure 19. Weight on Recession State, PSUR Model

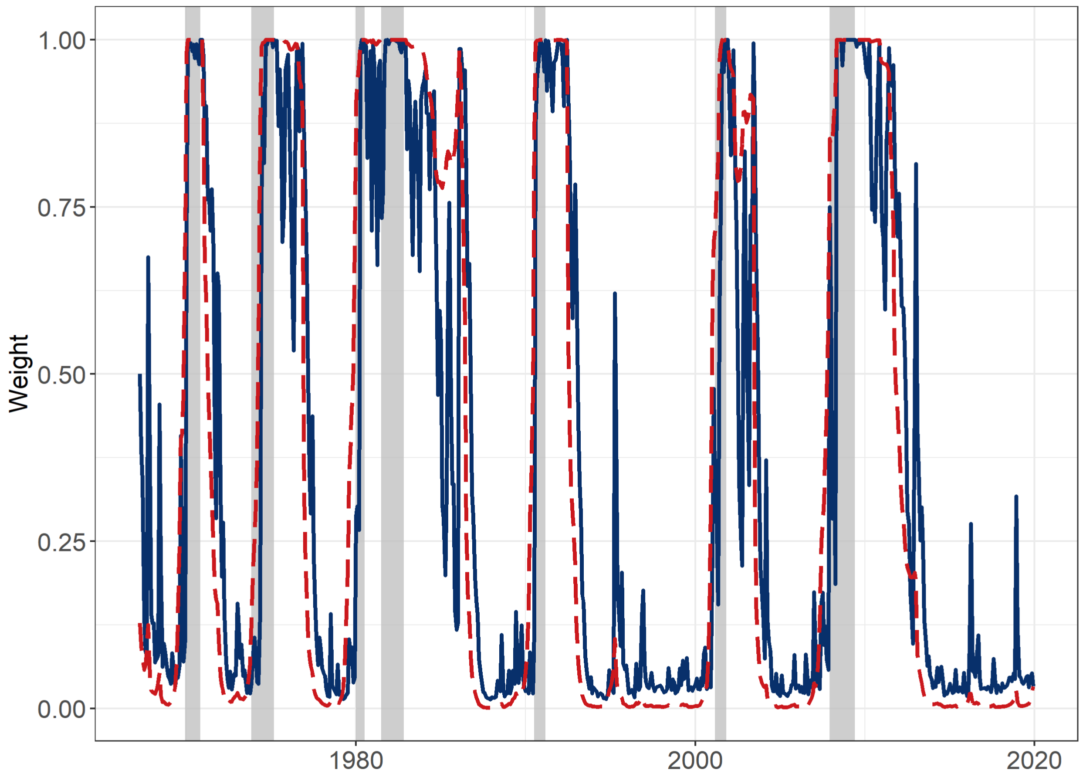

Note: The solid blue line is the filtered probability of being in a recession state at time t; the dashed red line is the smoothed probability of being in a recession state at time t. Gray bars denote periods identified as recessions by the National Bureau of Economic Research. PSUR = permanently separated unemployment rate.

These estimated probabilities coincide very closely with the probabilities estimated using the quarterly unemployment-gap data and with recessions as defined by NBER.

For the TSUR, I use a standard time-series model, fitting an AR(4) with monthly TSUR data over the sample spanning May 1967 through December 2019. Four lags are chosen because this value minimizes the Bayes-Schwarz information criterion.[33] I exclude 2020 data from the estimation sample for the TSUR given the unprecedented changes since the start of the pandemic. Although I restrict the estimation sample to prepandemic data for both the PSUR and TSUR models, the estimated dynamics should still produce a reasonable forecast

[33]The Akaike information criterion is minimized at higher-order lags, but I use four lags for a more parsimonious model.

conditional on recent data, without the drawback of significantly impacting the estimated parameters.

6.2 Simulations

The procedure for constructing simulations for each model is similar to the method described in Section 3.5. I create simulations of the PSUR and TSUR using the two models discussed above. Shocks for the PSUR are drawn from the sample of estimated errors of the PSUR model, with the probability of drawing shock $\hat{\varepsilon}_t$ in state i proportional to the probability that the economy was in state i in period t. Shocks for the TSUR are drawn from the sample of estimated errors of the TSUR model, with the drawn shock corresponding to the same time period as the PSUR shock. For the initial conditions of the PSUR and TSUR, I use data available as of the Bureau of Labor Statistics' November 2021 employment report. Additionally, all simulations start in the expansion state. This initial value is chosen to reflect the fact that the PSUR was declining, on average, in the last several months of the available data. After the initial simulation period, simulations transition endogenously between states with probabilities determined using the estimated transition probability parameters.

These monthly simulations of the TSUR and PSUR are created over the December 2021–March 2022 horizon. The monthly simulations are then converted to quarterly series of the unemployment rate using averages across each quarter. I then convert these unemployment-rate simulations to unemployment-gap simulations by subtracting CBO's estimates of the noncyclical unemployment rate. These quarterly values are then used as initial conditions for simulations created using the quarterly Markov-switching model of the unemployment gap. This two-step procedure allows me to properly account for the unusual split between temporarily and permanently separated unemployment in the near term while using the benchmark Markov-switching model of the unemployment gap to create simulations of the unemployment rate over longer forecast horizons.

One final change is made to the quarterly model of the unemployment gap. I adjust the constant term in the equation governing the probability of staying in the recession state

Figure 20. Fan Chart of Unemployment-Rate Simulations

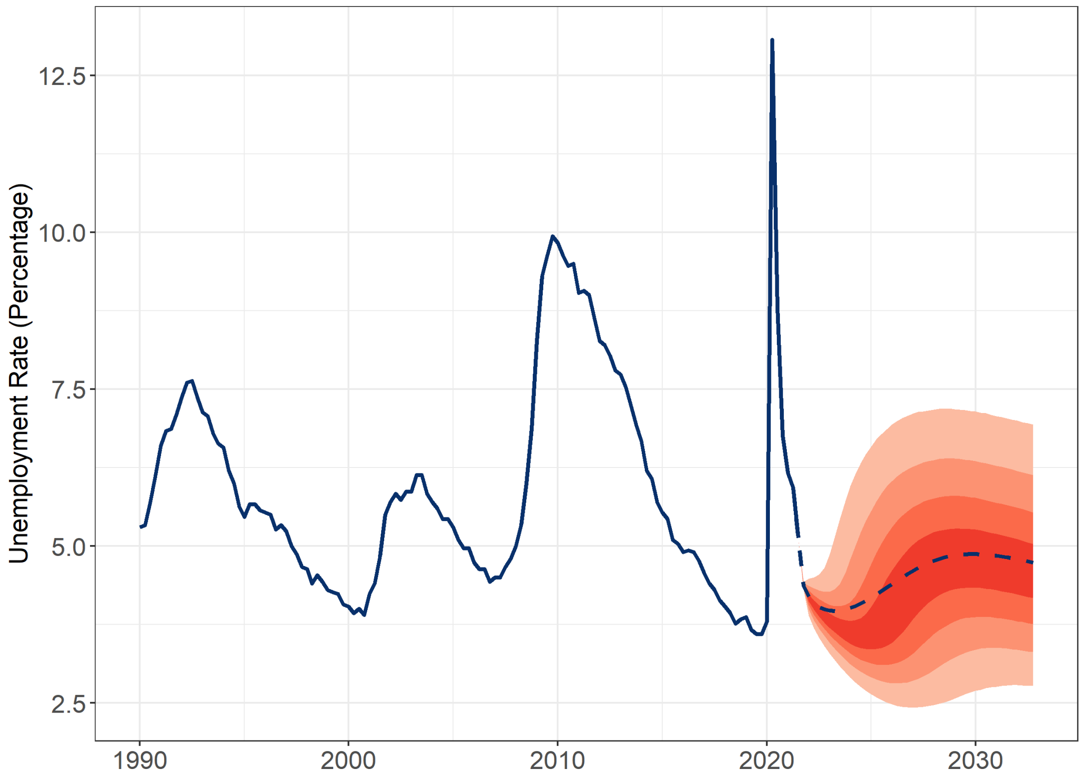

Note: The dashed blue line is the average across simulations in each time period, with simulations starting in Quarter 4 (Q4) of 2021. Shaded red intervals are quantiles of the simulations, from 10 percent to 90 percent in the largest range (lightest red) and from 40 percent to 60 percent in the smallest (darkest red).

from 2.42 to 1.90. For reference, the standard error on the estimated parameter is 0.485, suggesting this change is minor relative to the uncertainty around the point estimate of the parameter. This change has the effect of making the long-run average unemployment gap across simulations consistent with CBO's long-run forecast of the unemployment gap.

I produce 100,000 simulations of the unemployment gap using the methodology described above. I then convert simulations of the unemployment gap to simulations of the unemployment rate by adding CBO's estimate of the noncyclical rate of unemployment (which I assume is exogenous) in each quarter of the forecast horizon. Figure 20 displays a fan chart of the unemployment rate over the forecast horizon. The dashed line in the figure is the

Figure 21. Five Example Simulations From Markov-Switching Model

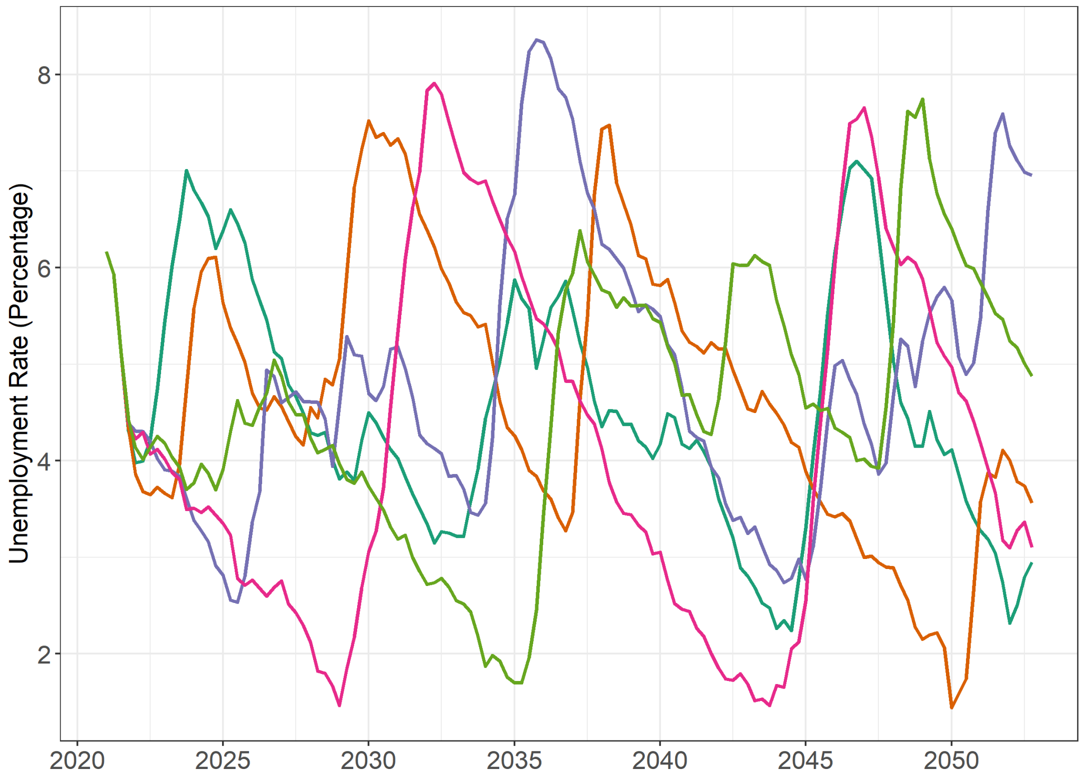

Note: The graph plots simulations of the unemployment rate from the Markov-switching model of the unemployment gap with time-varying transition probabilities.

average unemployment rate across simulations in each quarter, and the shaded intervals represent quantiles of the simulations in each time period (from 10 percent to 90 percent in the largest range and from 40 percent to 60 percent in the smallest). As shown in the figure, most simulations of the unemployment rate gradually decline over the next several years, with the mean across simulations gradually stabilizing. However, the dispersion across simulations increases significantly over time as the simulations transition between expansion and recession. To illustrate the dynamics of individual simulations, Figure 21 displays five example simulations over a 30-year forecast horizon. Each simulation exhibits periodic recessions, which are characterized by a rapidly increasing unemployment rate followed by more gradual declines. This is the asymmetric dynamic that I was trying to capture and, as shown in Section 3.5, these simulations can be characterized as "recessionary" at a rate consistent with the historical data, whereas simulations produced by the linear model cannot.

7 Conclusion

In this paper I have reviewed the evidence that the unemployment rate exhibits asymmetry over the business cycle. I have also shown that a Markov-switching model estimated on the historical data captures this observed asymmetry, with the unemployment rate rising faster in recession than it falls in expansion. Moreover, this result is invariant to both the de-trending of the unemployment rate and the particular assumption of transition probabilities.

This model also produces recession simulations at a rate consistent with the historical data, whereas a simpler linear version of the model does not. I also showed that a statistical test of duration dependence in the business cycle has almost no power to detect this feature when the sample size of observations is small and the effect of duration dependence is indirect. Additionally, I presented evidence that the benchmark Markov-switching model produces forecasts superior to those of a simpler version of the model with CTPs, in addition to the linear version of the model.

To account for the unique split between permanently separated and temporarily separated unemployed workers in the pandemic recession, I estimated a version of the Markov-switching model on the permanently separated unemployment rate. This allowed me to produce simulations that reflected the unique dynamics observed during the recent recession and recovery.

Future work by CBO could detail how this model and simulations are used for estimating the cost of relevant programs in the agency's baseline forecast, as well as how these simulations are used to estimate the cost of relevant legislative proposals.

A Estimation

The model is estimated using an expectation-maximization (EM) algorithm, which proceeds in two steps. In the first (expectation) step, state probabilities are formed conditional on the observed data and an estimate of the parameter vector. In the second (maximization) step, a new estimate of the parameter vector is formed conditional on the observed data and state probabilities from the expectation step. The procedure is then repeated, using the estimated parameter vector from the maximization step in the previous iteration in the expectation step in the current iteration. This process is repeated until the log likelihood of the sample process converges. Let $\theta^{(j)} = (\beta_1^{(j)}, \beta_2^{(j)}, \sigma_1^{(j)}, \sigma_2^{(j)}, \zeta_1^{(j)}, \zeta_2^{(j)})$ be the vector of parameters estimated in the jth iteration of the algorithm.

A.1 Expectation Step

In the expectation step, I form state probabilities using the estimated vector of parameters from the maximization step in the previous iteration. Let $\mathcal{F}_t = (\widetilde{u}_t, \widetilde{u}_{t-1}, ..., \widetilde{u}_1)$ be a vector containing all observations through date t. Using the fact that $\varepsilon_t \sim N(0,1)$, I can write the density of $\widetilde{u}_t$ conditional on being in state i, past observations, and the estimated parameters from the maximization step in the last iteration as

$$f(\widetilde{u}_t | s_t = i, \mathcal{F}_{t-1}; \beta_i^{(j-1)}, \sigma_i^{(j-1)}) = \frac{1}{\sqrt{2\pi}\sigma_i^{(j-1)}} \exp\left\{\frac{-(\widetilde{u}_t - X_t'\beta_i^{(j-1)})^2}{2\sigma_i^{(j-1)}}\right\} \tag{3}$$

Noting that there is a density for each state, let η_t be a (2×1) vector of the conditional densities at time t. Additionally, let $P(s_t = i | \mathcal{F}_t; \beta_i^{(j-1)}, \sigma_i^{(j-1)})$ be the probability of being in state i at time t conditional on all observations up to and including time t and the estimated parameters from the last iteration. Then I can collect these probabilities for each state into a (2×1) vector $\xi_{t|t}$. Similarly, let $P(s_t = i | \mathcal{F}_{t-1}; \beta_i^{(j-1)}, \sigma_i^{(j-1)})$ be the probability of being in state i at time t conditional on all observations up to and including time $t-1$ and the estimated parameters from the last iteration, and collect these probabilities for each state into a (2×1) vector $\xi_{t|t-1}$. Then the optimal inference for each date t in the sample is given by

$$\xi_{t|t} = \frac{(\xi_{t|t-1} \odot \eta_t)}{1'(\xi_{t|t-1} \odot \eta_t)} \tag{4}$$

where $\odot$ denotes element-by-element multiplication and $1'$ is a (1×2) vector of ones.

In order to estimate $\xi_{t|t}$ for each period t, I need estimates of $\xi_{t|t-1}$, the vector of state probabilities for period t given observations up to time $t-1$. For this, I assume the state of the economy, s_t, evolves according to a Markov chain, with time-varying transition probabilities. The probability of transitioning from state i in period $t-1$ to state i in period t, conditional on past observations and the estimated parameters from the last iteration, is given by

$$P(s_t = i | s_{t-1} = i, \mathcal{F}_{t-1}; \zeta_i^{(j-1)}) = \frac{\exp(z_t' \zeta_i^{(j-1)})}{1 + \exp(z_t' \zeta_i^{(j-1)})} = p_{i,t} \tag{5}$$

Let P_t be the transition probability matrix at time t, with $p_{ij,t}$ being the element in the ith row and jth column. Then the vector of state probabilities for period t given observations up to time $t-1$ is given by

$$\xi_{t|t-1} = P_t' \xi_{t-1|t-1} \tag{6}$$

Thus, given an initial value $\xi_{0|0}$ and the estimated parameters from the last iteration, I can iterate on Equations 4 and 6, using the transition probability matrix from Equation 5, to calculate $\xi_{t|t}$ for each time t.

I would also like to estimate probabilities of being in a given state at time t using observations not only up to time t but through the end of the sample T. Kim (1994) proposed an algorithm for estimating the probability of being in each state at time t using observations up through time $T > t$. The algorithm works by calculating

$$\xi_{t|T} = \xi_{t|t} \odot \{P_{t+1}[\xi_{t+1|T}(\div)\xi_{t+1|t}]\} \tag{7}$$

starting with $t = T - 1$ and iterating back to $t = 1$, where $(\div)$ denotes element-by-element division.

This completes the expectation step, and I will use the state probabilities formed in Equation 7 in the maximization step.

A.2 Maximization Step

In the maximization step, I estimate the model parameters using the state probabilities formed in the expectation step in the current iteration. Starting with the parameters governing transition probabilities, Diebold et al. (1994) showed that maximum likelihood estimates for these parameters satisfy

$$\begin{aligned}\zeta_i^{(j)} =& \left[\sum_{t=2}^{T} z_t P(s_{t-1} = i|\mathcal{F}_T; \theta^{(j-1)}) \frac{\partial p_{i,t}}{\partial \zeta_i}\right]^{-1} \\ &\cdot \left[\sum_{t=2}^{T} z_t \Big\{P(s_t = i, s_{t-1} = i|\mathcal{F}_T; \theta^{(j-1)}) - P(s_{t-1} = i|\mathcal{F}_T; \theta^{(j-1)}) \Big(p_{i,t} - \frac{\partial p_{i,t}}{\partial \zeta_i} \zeta_i^{(j-1)}\Big)\Big\}\right]\end{aligned} \tag{8}$$

where the partial derivatives are evaluated at $\zeta_i^{(j-1)}$. This equation comes from a first-order Taylor series expansion of $p_{i,t}$ around $\zeta_i^{(j-1)}$, which gives me a closed-form solution for ζ_i.

For β_i, Hamilton (1994) showed that these parameters satisfy an orthogonality condition with observations weighted by the probability of being in each state. The probabilities used to weight the observations are the probabilities formed in Equation 7 in the expectation step.

Letting

$$\bar{u}_t(i) = \widetilde{u}_t \cdot \sqrt{P(s_t = i|\mathcal{F}_T; \theta^{(j-1)})} \tag{9}$$

$$\bar{X}_t(i) = X_t \cdot \sqrt{P(s_t = i|\mathcal{F}_T; \theta^{(j-1)})} \tag{10}$$

be the weighted observations, then the ordinary least squares (OLS) estimates of β_i are given by

$$\beta_i^{(j)} = \left[\sum_{t=2}^{T} [\bar{X}_t(i)][\bar{X}_t(i)]'\right]^{-1} \left[\sum_{t=2}^{T} [\bar{X}_t(i)]\bar{u}_t(i)\right] \tag{11}$$

Finally, the estimates of σ_i are given by

$$\sigma_i^{(j)} = \left(\frac{\sum_{t=2}^{T} (\widetilde{u}_t - X_t'\beta_i^{(j)})^2 \cdot P(s_t = i|\mathcal{F}_T; \theta^{(j-1)})}{\sum_{t=2}^{T} P(s_t = i|\mathcal{F}_T; \theta^{(j-1)})}\right)^{1/2} \tag{12}$$

This completes the maximization step. With iteration j complete, I now run the expectation step for iteration $j+1$, using the parameters estimated in the maximization step in iteration j.

A.3 Starting the Algorithm

To start the algorithm, I need values for the initial state probabilities, $\xi_{1|0}$, and initial values for the parameter vector, $\theta^{(0)}$. For $\xi_{0|0}$, I use uniform probabilities, attaching equal initial probabilities to each state. $\theta^{(0)}$ is chosen using a random search, selecting the parameters associated with the highest log likelihood. The random values for $\beta_i^{(0)}$ are drawn from a normal distribution with mean equal to the OLS estimate of the parameter in the linear model and with standard deviation equal to the Newey-West (1987) standard error. Random values for $\sigma_i^{(0)}$ are drawn from a normal distribution with mean equal to the standard deviation of the fitted error from the linear model estimated with OLS and with standard deviation equal to 0.1 times the standard deviation of the fitted error from the linear model estimated with OLS. Random values for $\zeta_i^{(0)}$ are drawn from the standard normal distribution. One thousand vectors of θ are drawn using this method, and the parameter vector with the highest log likelihood is chosen as $\theta^{(0)}$.

References

[1] Mirko Abbritti and Stephan Fahr. Downward Wage Rigidity and Business Cycle Asymmetries. *Journal of Monetary Economics*, 60(7):871–886, 2013.

[2] David Andolfatto. Evidence and Theory on the Cyclical Asymmetry in Unemployment Rate Fluctuations. *The Canadian Journal of Economics*, 30(3):709–721, 1997.

[3] Jushan Bai and Serena Ng. Tests for Skewness, Kurtosis, and Normality for Time Series Data. *Journal of Business and Economic Statistics*, 23(1):49–60, 2005.

[4] Arthur F. Burns and Wesley C. Mitchell. *Measuring Business Cycles.* National Bureau of Economic Research, New York, 1946.

[5] J. Bradford DeLong and Lawrence H. Summers. Are Business Cycles Symmetrical? In Robert J. Gordon, editor, *The American Business Cycle: Continuity and Change*, pages 166–178. University of Chicago Press for the National Bureau of Economic Research, Chicago, IL, 1986.

[6] Wouter J. Den Haan, Garey Ramey, and Joel Watson. Job Destruction and Propagation of Shocks. *American Economic Review*, 90(3):482–498, 2000.

[7] Peter A. Diamond. Aggregate Demand Management in Search Equilibrium. *Journal of Political Economy*, 90(5):881–894, 1982.

[8] Francis X. Diebold, Joon-Haeng Lee, and Gretchen C. Weinbach. Regime Switching With Time-Varying Transition Probabilities. In *Nonstationary Time Series Analysis and Cointegration*, pages 283–302. 1994.

[9] Francis X. Diebold and Roberto S. Mariano. Comparing Predictive Accuracy. *Journal of Business & Economic Statics*, 13(3):253–263, 1995.

[10] Francis X. Diebold and Glenn D. Rudebusch. A Nonparametric Investigation of Duration Dependence in the American Business Cycle. *Journal of Political Economy*, 98(3):596–616, 1990.

[11] Stephane Dupraz, Emi Nakamura, and Jón Steinsson. *A Plucking Model of Business Cycles*. NBER Working Paper 26351, 2020.

[12] J. Michael Durland and Thomas H. McCurdy. Duration-Dependent Transitions in a Markov Model of U.S. GNP Growth. *Journal of Business and Economic Statistics*, 12(3):279–288, 1994.

[13] Barry Falk. Further Evidence on the Asymmetric Behavior of Economic Time Series Over the Business Cycle. *Journal of Political Economy*, 94(5):1096–1109, 1986.

[14] Domenico Ferraro. The Asymmetric Cyclical Behavior of the U.S. Labor Market. *Review of Economic Dynamics*, 30(October):145–162, 2018.

[15] Jessica Gallant, Kory Kroft, Fabian Lange, and Matthew Notowidigdo. Temporary Unemployment and Labor Market Dynamics During the COVID-19 Recession. *Brookings Papers on Economic Activity*, Fall, 2020.

[16] Raffaella Giacomini and Barbara Rossi. Forecast Comparisons in Unstable Environments. *Journal of Applied Econometrics*, 25(4):595–620, 2010.

[17] Raffaella Giacomini and Halbert White. Tests of Conditional Predictive Ability. *Econometrica*, 74(6):1545–1578, 2006.

[18] Manuel González-Astudillo and Diego Vilán. A New Procedure for Generating the Stochastic Simulations in FRB/US. *Finance and Economics Discussion Series Notes*, March 2019.

[19] Robert E. Hall and Marianna Kudlyak. *Why Has the US Economy Recovered So Consistently from Every Recession in the Past 70 Years?* NBER Working Paper 27234, 2020.

[20] James D. Hamilton. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57(2):357–384, 1989.

[21] James D. Hamilton. *Time Series Analysis*. Princeton University Press, Princeton, NJ, 1994.

[22] John Maynard Keynes. *The General Theory of Employment, Interest, and Money*. Macmillan, London, 1936.

[23] Chang Jin Kim. Dynamic Linear Models With Markov-Switching. *Journal of Econometrics*, 60(1-2):1–22, 1994.

[24] Chang-jin Kim and Charles R. Nelson. Business Cycle Turning Points, a New Coincident Index, and Tests of Duration Dependence Based on a Dynamic Factor Model With Regime-Switching. *The Review of Economics and Statistics*, 80(2):188–201, 1998.

[25] Britta Kohlbrecher, Christian Merkl, and Daniela Nordmeier. Revisiting the Matching Function. *Journal of Economic Dynamics and Control*, 69(August):350–374, 2016.

[26] Alisdair McKay and Ricardo Reis. The Brevity and Violence of Contractions and Expansions. *Journal of Monetary Economics*, 55(4):738–751, 2008.

[27] Dale T. Mortensen and Christopher A. Pissarides. Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies*, 61(3):397–415, 1994.

[28] Salih N. Neftci. Are Economic Time Series Asymmetric over the Business Cycle? *Journal of Political Economy*, 92(2):307–328, 1984.

[29] Whitney K. Newey and Kenneth D. West. A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703–708, 1987.

[30] Christopher Pissarides. Search Unemployment With On-the-Job Search. *Review of Economic Studies*, 61(3):457–475, 1994.

[31] Philip Rothman. Further Evidence on the Aymmetric Behavior of Unemployment Rates Over the Business Cycle. *Journal of Macroeconomics*, 13(2):291–298, 1991.

[32] Claudia Sahm. Direct Stimulus Payments to Individuals. In *Recession Ready: Fiscal Policies to Stabilize the American Economy*, pages 67–92. The Hamilton Project and the Washington Center on Equitable Growth, Washington, DC, 2019.

[33] S. S. Shapiro and M. B. Wilk. An Analysis of Variance Test for the Exponential Distribution (Complete Samples). *Technometrics*, 14(2):355–370, 1972.

[34] Daniel E. Sichel. Are Business Cycles Asymmetric ? A Correction. *Journal of Political Economy*, 97(5):1255–1260, 1989.

[35] Daniel E. Sichel. Business Cycle Duration Dependence: A Parametric Approach. *The Review of Economics and Statistics*, 73(2):254–260, 1991.

[36] Daniel E. Sichel. Business Cycle Asymmetry: A Deeper Look. *Economic Inquiry*, 31(2):224–236, 1993.

[37] Kenneth D. West. Asymptotic Inference About Predictive Ability. *Econometrica*, 64(5):1067–1084, 1996.

[38] Halbert White. Maximum Likelihood Estimation of Misspecified Models. *Econometrica*, 50(1):1–25, 1982.

[39] Thomas W. Zuehlke. Business Cycle Duration Dependence Reconsidered. *Journal of Business and Economic Statistics*, 21(4):564–569, 2003.