

WHATEVER HAPPENED TO THE BUSINESS CYCLE? A BAYESIAN ANALYSIS OF JOBLESS RECOVERIES

KRISTIE M. ENGEMANN AND MICHAEL T. OWYANG

Federal Reserve Bank of St. Louis

During the typical recovery from U.S. postwar period economic downturns, employment recovers to its prerecession level within months of the output trough. However, during the past two recoveries, employment has taken up to three years to achieve its prerecession benchmark. We propose a formal empirical model of business cycles with recovery periods to demonstrate that the past two recoveries have been statistically different from previous experiences. We find that this difference can be attributed to a shift in the speed of transition between business cycle regimes. Moreover, we find this shift results from both durable and nondurable manufacturing sectors losing their cyclical characteristics. We argue that this finding of acyclicity in post-1980 manufacturing sectors is consistent with previous hypotheses (e.g., improved inventory management) regarding the reduction in macroeconomic volatility over the same period. These results suggest a link between the two phenomena, which have heretofore been studied separately.

Keywords: Business Cycles, Gibbs Sampling, Smooth Transition Autoregression, Industrial Disaggregation

1. INTRODUCTION

The notion that the economy moves between distinct, alternating phases is often attributed to the pioneering work of Burns and Mitchell (1946). They argue that economic dynamics change between expansionary and recessionary phases since dubbed the business cycle. More recently, Hamilton (1989) developed a statistical representation of this idea by assuming that the transitions between the two phases of the economy are governed by a two-state Markov process. Hamilton's assumption was that the current state of the economy was determined by a probability that depended only on the past period's phase. A further innovation consistent with the idea of the economy residing in distinct phases was introduced by Kim et al. (2005), in which they argue that the postwar U.S. economy has experienced

The authors benefitted from conversations with Kathryn Birkeland, Mike Dueker, Jason Faberman, James Morley, and Jeremy Piger. We acknowledge comments from participants at the 2005 WEAI, the 2007 SNDE, the 2007 CEF, and seminar participants at the University of California, Davis; the University of Texas, Dallas; and West Virginia University. Heidi Beyer-Powe provided research assistance. The views expressed here are the authors' and do not necessarily reflect those of the Federal Reserve Bank of St. Louis or the Federal Reserve System. All of the usual disclaimers apply. Address correspondence to: Michael T. Owyang, Research Division, Federal Reserve Bank of St. Louis, P.O. Box 442, St. Louis, MO 63166, USA; e-mail: owyang@stls.frb.org.

a rapid recovery following each recession phase. This recovery period—which they term the *bounce back*—typically has lasted for six quarters subsequent to the business cycle trough and is characterized by higher-than-average GDP growth.

For most postwar recoveries, the timing of the return to prerecession employment levels has lagged the GDP turning point by only a few months. However, the experiences of the U.S. economy following the two most recent recessions have been decidedly different from previous recoveries. During the periods subsequent to the troughs of the 1990–1991 and 2001 recessions, employment growth was not as strong as it was in previous periods. In fact, for many months after the NBER-determined turning points, employment growth was negligible or even negative. Thus, employment took many more months to return to prerecession levels. This has led both the media and academics to term these periods the *jobless recoveries*.

Although differences between these jobless recoveries and previous recoveries have been well documented, the cause of the change has been subject to interpretation. One possible reason for the variety of interpretations is that the evidence for jobless recoveries has not been formalized in rigorous statistical models. For example, there is still some controversy about the number of jobless recoveries and whether trend employment or prerecession employment is the proper measure of the point of recovery.¹ These controversies may result from the fact that recoveries account for such a small proportion of the postwar period. Because only ten such periods exist, it remains possible that each recovery exhibits unique anecdotal and statistical properties.²

Evaluating such a claim necessitates the construction of a formal empirical model, which we propose in this paper. In particular, we appeal to the literature on nonlinear time series representations of the business cycle à la Hamilton (1989) and others. Our model of choice is the smooth transition autoregression [Teräsvirta and Granger (1993)], which allows the change in regime to occur slowly over time. This change in regime captures the spirit of the empirical business cycle models in that the steady-state employment growth rate and the dynamics differ across the two regimes. However, a shift into a different business cycle phase does not occur instantaneously but is governed by a speed-of-adjustment parameter. The variation in speed of adjustment allows us to explicitly model the slow recoveries witnessed in the early 1990s and 2000s and determine, counterfactually, the resulting job loss over the previous typical recessions of the same duration.

The balance of the paper is organized as follows. Section 2 reviews the evidence for jobless recoveries in the two most recent recessions. Section 3 presents the baseline smooth transition autoregressive model, our formal representation of jobless recoveries, which nests the two-phase threshold model.³ Section 4 outlines the algorithm used to estimate the model. Section 5 presents the results from the estimation of the baseline smooth transition model, allowing for a structural break in the rate of transition between recession and expansion. In addition, we estimate the magnitude of aggregate job loss caused by jobless recoveries by conducting




FIGURE 1. Aggregate payroll employment.

counterfactual experiments. Section 6 considers the causes of jobless recoveries by reestimating the model with industrially disaggregated data. Section 7 concludes.

2. REEXAMINING THE EVIDENCE FOR JOBLESS RECOVERIES

The existence of jobless recoveries has primarily been documented as an increase in the number of months that the employment trough lags the GDP trough. Figure 1 plots employment, indexed to its prerecession level, following the past two recessions and the means and extrema for the other seven recessions dating back to 1953.⁴ This figure illustrates two features inherent in the jobless recoveries. First, after ten months, employment relative to the prerecession peak is uniformly below the minimum for all previous nonjobless recoveries. Second, the lag between the NBER turning point and the resumption of positive employment growth for the jobless episodes is substantially longer.

Although there seems to be consensus regarding the stylized facts, the causes of jobless recoveries are subject to interpretation.⁵ Four hypotheses have been suggested: (i) organizational restructuring (i.e., an increase in the severance rate), (ii) sectoral reallocation, (iii) innovations in labor demand, and (iv) compositional changes in labor supply.

Proponents of the organizational restructuring hypothesis suggest that a relationship exists between the duration of an expansion and the duration of recovery. Koenders (2005) finds a positive correlation between the time by which employment recovery lags output recovery and longer expansions, an observation that holds across all sectors for the jobless recoveries. In a related paper, Koenders and Rogerson (2005) explain this link by developing an organizational restructuring model where firms wait to eliminate labor inefficiencies until demand for

production declines and postpone hiring during periods of reorganization. Thus, longer expansions imply that firms had more inefficiencies to eliminate during the recession period.⁶

Advocating structural change across sectors as a cause for jobless recovery, Groshen and Potter (2003) categorize industries based on whether job adjustment is cyclical (i.e., temporary) or structural (i.e., permanent). They show that the proportion of employment in structurally changing industries jumped from 57% in 1990–1991 to 79% for the 2001 recession. Because employers find it more difficult to create new positions than simply to recall workers, Groshen and Potter argue that the increase in permanent job changes has contributed to the joblessness of the current recovery.⁷

A third explanation hinges on an unusual fall in labor demand and/or a hiring slowdown. Schreft and Singh (2003) find that the substitution of just-in-time employment (in the form of temporary and part-time workers) for more permanent employment remains higher in the two jobless recoveries than in previous recoveries. By having the option of more flexible employment, firms can lay off workers but can call on just-in-time employment if necessary.⁸ A fourth explanation centers on recent changes in labor supply and argues that the employment situation might not be as bad as it seems. Aaronson et al. (2004a) conclude that there has been a slight increase in the level of unincorporated self-employed workers, who are not included in payroll employment numbers, since the start of the 2001 recession. They argue that unincorporated employment might have increased because of the increased difficulty of finding jobs.⁹

3. AN EMPIRICAL MODEL OF THE BUSINESS CYCLE

There are a number of methods for developing time series models of business cycles and, hence, recovery periods. However, our objective is to model jointly the distinct differences in dynamics between the recession and expansion phases in the economy and the slow transition between these two phases exhibited in the past two recessions.¹⁰ A vast literature has shown that Markov-switching models are one method of elucidating the varying dynamics occurring between business cycle phases.¹¹ However, the Markov-switching model does not explicitly allow the gradual transition dynamics necessary to model the recoveries. Our intention is to parsimoniously unify the literature on jobless recoveries and empirical business cycle models. Thus, we employ the smooth transition autoregressive (STAR) model as our model of the business cycle.

Conditional on a predetermined path for the threshold variable, the simple STAR model does not exhibit significant differences in transition dynamics across recessions. In other words, the model is not sufficiently rich to generate shifts in the recovery dynamics without deriving these from differences in the transition variable. Thus, we incorporate simultaneous structural breaks in both the transition function parameters and the steady-state employment growth rates.¹²

The STAR(p) model is essentially a p -lag autoregression of the form

$$y_t = \alpha_0 + \alpha_1 \pi(z_{t-d}) + \sum_{i=1}^p [\phi_{0i} + \phi_{1i} \pi(z_{t-d})] y_{t-i} + \varepsilon_t, \tag{1}$$

where y_t is the period- t variable of interest, $\theta = [\alpha_0, \alpha_1, \phi_{01}, \phi_{11}, \dots, \phi_{0p}, \phi_{1p}]'$ is a vector of parameters, z_{t-d} is the d th lag of a variable that governs the state of the economy, and the error term $\varepsilon_t \sim N(0, \sigma^2)$. The difference between (1) and a simple autoregression lies in the transition function $\pi(z_{t-d}) \in [0, 1]$, which governs both the lag coefficients and the intercept term.

The path of the economy is uniquely determined by $\pi(z_{t-d})$, which is bounded between zero and one. Although the transition function can take a number of forms, we assume the logistic representation¹³

$$\pi(z_{t-d}; \gamma, c) = \{1 + \exp[-\gamma(z_{t-d} - c)]\}^{-1}, \tag{2}$$

where γ and c are parameters. Here, γ represents the transition speed of the economy between the extreme business cycle phases, represented by $\pi(z_{t-d}; \gamma, c) = 0$ and $\pi(z_{t-d}; \gamma, c) = 1$. As $\gamma \rightarrow \infty$, the model becomes a standard two-phase threshold autoregressive model. Conversely, as $\gamma \rightarrow 0$, the model approaches an AR(p). The parameter c governs the point at which the effect of the threshold variable changes sign. Finally, the delay parameter, d , determines the lag by which the transition variable affects the variable of interest.

This model captures business cycles in a manner similar to the Markov-switching model of Hamilton (1989), where the transition function $\pi(z_{t-d}; \gamma, c)$ is the analogue of the hidden Markov variable. During recessions, $\pi(z_{t-d}; \gamma, c) = 0$, implying that the steady-state employment growth rate is $\alpha_0[1 - \sum_{i=1}^p \phi_{0i}]^{-1}$. Conversely, during expansions, $\pi(z_{t-d}; \gamma, c) = 1$ and steady-state employment growth is $(\alpha_0 + \alpha_1)[1 - \sum_{i=1}^p (\phi_{0i} + \phi_{1i})]^{-1}$. Finally, for periods in which $1 > \pi(z_{t-d}; \gamma, c) > 0$, the economy transitions between these two regimes. The transition function depends on the threshold variable, z_{t-d} , which influences the state of the economy. For example, z_t might be a contemporaneous business cycle indicator such as GDP. When z_{t-d} exceeds the threshold c (say $c = 0$), the economy tends to move toward expansion.

4. ECONOMETRIC IMPLEMENTATION

The model in the preceding subsection can be estimated either via classical [e.g., Teräsvirta and Granger (1993)] or Bayesian [Lopes and Salazar (2006)] methods. The latter method, which allows computation of marginal likelihoods and Bayes factors, readily facilitates testing for changes in the structural parameters. Further, Bayesian methods can easily be extended to estimate the timing of the break date.

To account for a possible break in the model parameters and to facilitate estimation, we can rewrite (1) in the form

$$y_t = \theta_j x_{t,j} + \varepsilon_t, \tag{3}$$

where

$$\theta_j = [\alpha_{0,j}, \alpha_{1,j}, \phi_{01,j}, \phi_{11,j}, \dots, \phi_{0p,j}, \phi_{1p,j}],$$

$$x_{t,j} = [1, \pi_j(z_{t-d}), y_{t-1}, \pi_j(z_{t-d})y_{t-1}, \dots, y_{t-p}, \pi_j(z_{t-d})y_{t-p}]'$$

$\varepsilon_t \sim N(0, \sigma_j^2)$, and $j = \{0, 1\}$ is an indicator variable that denotes the structural break τ —that is, if $t < \tau$, $j = 0$; conversely, if $t \geq \tau$, $j = 1$. Estimation of the posterior distributions of the model parameters can be accomplished via Gibbs sampling [Gelfand and Smith (1990)]. Suppose (θ_j, σ_j^2) have prior distributions $\theta_j \sim N(\bar{\theta}, \sigma_\theta^2 I_{2+2p})$ and $\sigma_j^2 \sim IG(v, \eta)$, respectively, where $IG(\cdot, \cdot)$ is the inverse gamma distribution. Then, conditional on $\pi_j(z_{t-d})$, drawing from the posterior distributions for the parameters of (3) is straightforward.

The vector of parameters $\psi_j = [\gamma_j, c_j, d_j]'$ from (2), however, may not have standard analytical posterior distributions but can be estimated via a Metropolis–Hastings step within the Gibbs sampler [Chib and Greenberg (1995)]. The Metropolis–Hastings algorithm requires a candidate drawn from a proposal density and an acceptance probability computed from the marginal likelihoods. Suppose that γ_j and c_j have prior distributions $\gamma \sim G(a, b)$ and $c \sim N(\bar{c}, \sigma_c^2)$, where $G(\cdot, \cdot)$ is the gamma distribution. Then a candidate $(\hat{\gamma}, \hat{c})$ can be drawn from the proposal densities

$$\hat{\gamma} \sim G\left(\frac{(\gamma^{[i]})^2}{\Delta_\gamma}, \frac{\gamma^{[i]}}{\Delta_\gamma}\right),$$

$$\hat{c} \sim N(c^{[i]}, \Delta_c),$$

where superscript i indicates the draw from the i th iteration and Δ_γ and Δ_c are chosen hyperparameters. The candidate is then accepted with probability $p = \min\{A, 1\}$, where

$$A = \frac{\prod_t dN(y_t | \pi(z_{t-d} | \hat{\gamma}, \hat{c}), \theta, \sigma^2)}{\prod_t dN(y_t | \pi(z_{t-d} | \gamma_0^{[i]}, c_0^{[i]}), \theta, \sigma^2)}$$

$$\times \frac{dG(\hat{\gamma} | a, b) dN(\hat{c} | \bar{c}, \sigma_c^2)}{dG(\gamma^{[i]} | a, b) dN(c^{[i]} | \bar{c}, \sigma_c^2)} \frac{dG\left(\gamma^{[i]} \left| \frac{(\hat{\gamma})^2}{\Delta_\gamma}, \frac{\hat{\gamma}}{\Delta_\gamma}\right.\right)}{dG\left(\gamma^{[i]} \left| \frac{(\gamma^{[i]})^2}{\Delta_\gamma}, \frac{\gamma^{[i]}}{\Delta_\gamma}\right.\right)},$$

and $dN(\cdot)$ and $dG(\cdot)$ are normal and gamma probability densities, respectively [Lopes and Salazar (2006)]. The delay parameter, d , assuming a discrete uniform prior on $[1, d_{\max}]$ for each subsample, can be drawn from a posterior distribution in

TABLE 1. Priors for estimation

Parameter	Prior distribution	Hyperparameters
$[\alpha, \phi]'$	$N(\mathbf{m}_0, \sigma^2 \mathbf{M}_0)$	$\mathbf{m}_0 = [1; -1; \mathbf{I}_4]; \mathbf{M}_0 = \mathbf{I}_6$
σ^2	$\Gamma^{-1}\left(\frac{\nu_0}{2}, \frac{\delta_0}{2}\right)$	$\nu_0 = 1; \delta_0 = 1$
c	$N(\mathbf{c}_0, \mathbf{C}_0)$	$\mathbf{c}_0 = 0; \mathbf{C}_0 = 1$
γ	$\Gamma(\mathbf{g}_0, \mathbf{G}_0)$	$\mathbf{g}_0 = 1; \mathbf{G}_0 = 1$
d	$U(\mathbf{d}_0, \mathbf{d}_1)$	$\mathbf{d}_0 = 1; \mathbf{d}_1 = 4$
τ	$U(\tau_0, \tau_1)$	$\tau_0 = 1975 : 01; \tau_1 = 1990 : 12$

which the probability of drawing $d = \hat{d}$ is determined by the weighted likelihood associated with \hat{d} relative to all other possible values of d .¹⁴

The model accounts for the possibility of a single structural break in both the measurement and transition parameters, θ and ψ , at time τ . The break date, τ , can be estimated by adding a step to the Gibbs sampler [Carlin et al. (1992)]. Conditional on the draws for the parameter vectors, the break date can be drawn from the distribution

$$\tau \sim p(\tau | Y, Z, \theta_1, \theta_2, \psi_1, \psi_2) = \frac{L(Y; Z, \tau, \theta_1, \theta_2, \psi_1, \psi_2, \sigma_1^2, \sigma_2^2)}{\sum L(Y; Z, \tau, \theta_1, \theta_2, \psi_1, \psi_2, \sigma_1^2, \sigma_2^2)},$$

where

$$L(Y; Z, \tau, \theta_1, \theta_2, \psi_1, \psi_2, \sigma_1^2, \sigma_2^2) = \frac{\exp\left[-\sum_{i=1}^2 \frac{1}{2\sigma_i^2} \varepsilon_i^{(\tau)'} \varepsilon_i^{(\tau)}\right]}{\sigma_1^\tau \sigma_2^{T-\tau}},$$

$\varepsilon_1^{(\tau)}$ is the $\tau \times 1$ vector of prebreak errors conditional on τ , $\varepsilon_2^{(\tau)}$ is its $(T - \tau) \times 1$ postbreak counterpart, and we have implicitly assumed a discrete uniform prior on a subset of the sample period.¹⁵ Table 1 summarizes the prior distributions for the model parameters and shows the values for the prior hyperparameters used in the estimation. The Gibbs sampler is a sequential draw from each conditional posterior distribution, with convergence yielding the ergodic distribution for the entire parameter set, including the break date, conditional on the data. We discard the first 10,000 draws and save the following 10,000 draws to compute the joint posterior density. Finally, we can determine whether the model with a break is preferred over the model with no break by examining the Bayes factor.¹⁶

The data employed in the estimation are the annualized growth rates of nonagricultural payroll employment and household employment. Nonagricultural payroll employment is compiled from the Current Employment Statistics survey conducted by the Bureau of Labor Statistics (BLS). Household employment data are taken from the BLS's Current Population Survey. The aggregate sample spans 1962:01 to 2005:12, which covers a total of six postrecession recovery periods. The model is estimated at monthly frequency using the appropriate lagged employment growth as the threshold variable.¹⁷

TABLE 2. Posterior distributions for model parameters

Parameter/ model	Payroll employment growth			Household employment growth		
	Full	Prebreak	Postbreak	Full	Prebreak	Postbreak
α_0	-0.55 (-1.19, -0.14)	-0.94 (-1.75, -0.27)	-0.20 (-0.98, 0.42)	-0.20 (-0.56, -0.03)	-0.49 (-1.05, -0.12)	-0.01 (-0.84, 0.79)
α_1	1.75 (1.04, 2.83)	2.25 (1.33, 3.36)	1.75 (0.60, 3.17)	2.66 (2.24, 3.27)	3.52 (2.90, 4.22)	2.43 (1.05, 4.04)
ϕ_{01}	0.70 (0.57, 0.86)	0.71 (0.53, 0.89)	0.68 (0.31, 1.14)	0.01 (-0.11, 0.13)	0.24 (0.07, 0.41)	-0.23 (-0.61, 0.08)
ϕ_{11}	-0.53 (-0.74, -0.35)	-0.48 (-0.70, -0.26)	-0.70 (-1.61, -0.24)	-0.05 (-0.21, 0.11)	-0.30 (-0.50, -0.10)	0.22 (-0.29, 0.90)
ϕ_{02}	-0.18 (-0.39, 0.00)	-0.30 (-0.58, -0.03)	0.18 (-0.18, 0.69)	0.14 (0.04, 0.25)	0.12 (-0.01, 0.26)	0.05 (-0.34, 0.47)
ϕ_{12}	0.56 (0.37, 0.78)	0.66 (0.38, 0.94)	0.13 (-0.57, 0.56)	-0.05 (-0.20, 0.10)	-0.12 (-0.29, 0.06)	0.08 (-0.62, 0.82)
σ^2	4.13 (3.83, 4.48)	5.18 (4.59, 5.92)	3.01 (2.67, 3.41)	11.35 (10.52, 12.30)	12.81 (11.41, 14.38)	9.27 (8.10, 10.60)
γ	0.96 (0.57, 1.83)	1.53 (0.90, 1.99)	0.60 (0.10, 1.62)	1.10 (0.47, 1.87)	3.17 (2.21, 4.14)	0.20 (0.04, 0.94)
c	-0.05 (-0.45, 0.32)	-0.05 (-0.43, 0.35)	-0.04 (-0.58, 0.46)	0.02 (-0.31, 0.27)	-0.02 (-0.27, 0.33)	-0.05 (-0.52, 0.44)
d	2	2	2	3	3	3

Notes: Intervals for 80% coverage are shown in parentheses. Payroll employment data are taken to be the non-agricultural, annualized, seasonally adjusted growth rate compiled from the Current Employment Statistics Survey. Household employment data are also annualized, seasonally adjusted growth rates compiled from the Current Population Survey. The aggregate sample spans 1962:01 to 2005:12.

5. AGGREGATE EMPIRICAL RESULTS

Estimation of the model yields a joint posterior distribution for the parameter vectors θ and ψ and an imputed posterior distribution for the transition function for both the full-sample and single-break models. In addition, results for the break model include the posterior distribution for the break date, τ . Evidence from marginal likelihoods reveals that, in most cases, the single-break model is favored over the full-sample model.

The medians of the posterior distributions of the parameters and their associated 10% and 90% quantiles for both the full-sample and single-break models are illustrated in Table 2.¹⁸ Allowing for a single endogenously chosen structural break results in a median break date near the end of the 1980–1982 recession.¹⁹ In this case, both the steady-state employment growth rates and the conditional variance of employment differ across subsamples. The first row of Table 3 provides additional evidence of the break; the Bayes factor comparing the baseline model with a single break to the no-break alternative overwhelmingly favors the former.²⁰ These reductions in conditional variance are consistent with the notion of a volatility reduction occurring in both GDP [McConnell and Perez-Quiros (2000)] and employment [Owyang et al. (2008)]. Evidence of jobless recoveries, on the other hand, may be gleaned by examining the estimated transition function, $\pi(z_{t-d})$. Although the posterior distributions for the threshold parameter, c , indicate no statistically important change across subsamples, the rate of transition does

TABLE 3. Bayes factors for payroll employment

Model	Break v. no break	Break v. AR(2)	No break v. AR(2)
Aggregate	101.2	59.5	-41.6
Durables	157.0	149.0	-8.0
Nondurables	107.0	201.0	94.0
TTU	-129.3	82.3	211.6
Services	145.3	162.1	16.8

Notes: Bayes factors are computed as the ratios of the marginal likelihoods computed using methods described in Chib (1995) and Chib and Jeliazkov (2001) with a flat model prior. Per the Jeffreys scale, a negative BF indicates the second model is favored. BFs greater than 2.3 indicate overwhelming evidence in favor of the first model listed.

decline from the prebreak to the postbreak samples.²¹ The decline in the speed of transition across subsamples embodies the jobless recoveries. The median decline in the speed of transition is approximately 60% suggesting a longer period before payroll employment achieves a full expansionary state, $\pi(z_{t-d}) = 1$.

In addition to the posterior distributions for the transition parameters, Figure 2 illustrates the median of the posterior density for the transition function, $\pi(z_{t-d})$, for both payroll and household employment. The transitions between recession and expansion in the prebreak period are rapid, indicative of a high value of γ . Conversely, the transitions after the postbreak recessions are slower, indicating enduring recovery periods and reflecting a lower γ . When they are compared with the NBER turning points, the impact of the structural break becomes apparent. During the prebreak period, the increase in the transition function to full expansion lags the NBER turning point by only a few months. Subsequent to the break, the lag between the NBER turning point and employment recovery is substantially longer. Moreover, the break reflects a reduction in the cyclicity of the economy.

In addition to testing for shifts in parameters, we can assess the effect of the jobless recoveries on aggregate employment. In particular, we can generate posterior densities for the reduction in employment from a counterfactual level implied by imposing the postbreak business cycle characteristics ψ_1 in a model with prebreak steady state growth rates θ_0 . At each iteration of the Gibbs sampler, we initialize these Monte Carlo experiments with respective data from the 1990–1991 and 2001 recessions. We then compute the counterfactual path of employment conditional on the realized values for the threshold variable, ψ_1 , and θ_0 . The resulting set of draws of the counterfactual employment loss form the posterior distribution.

We compare actual payroll employment levels during the two recovery periods with the medians of the posterior densities for counterfactual employment levels in Figure 3. The tenth and ninetieth percentiles for the counterfactuals are also shown. The counterfactual employment loss three years after the trough is more than nine times larger for the 2001 recession than for the 1990–1991 recession. The total loss is about 180,000 for 1990–1991 and 1.69 million for 2001.

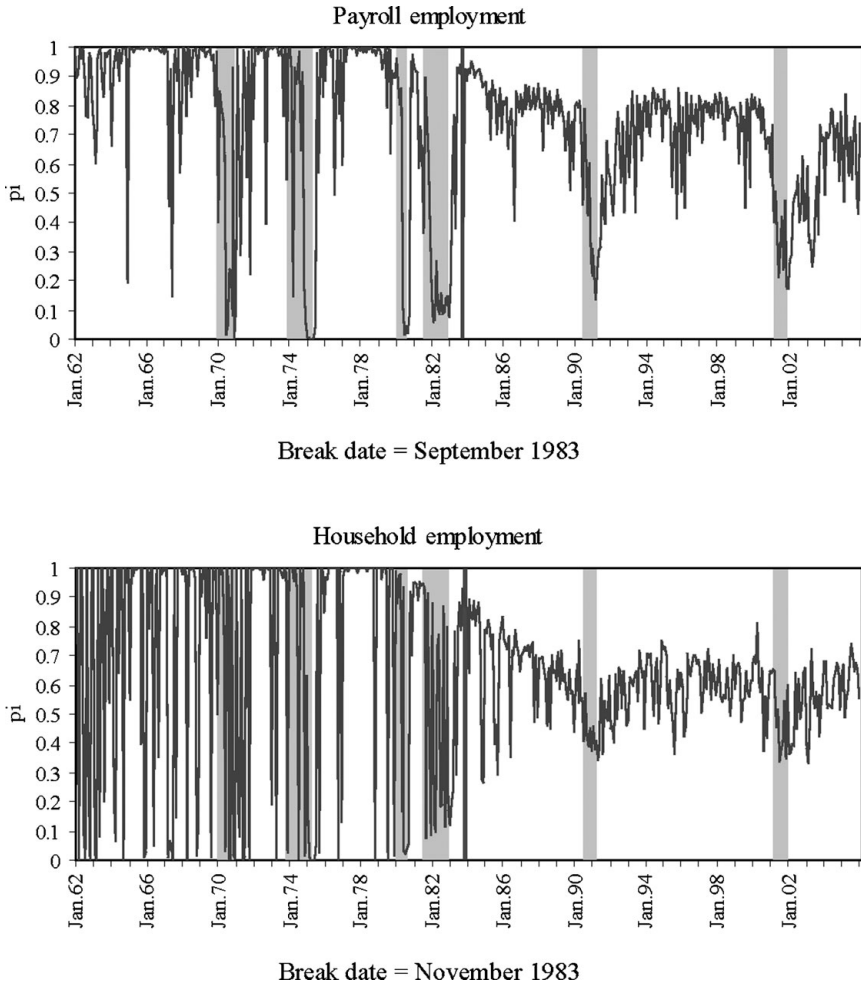


FIGURE 2. Transition functions with aggregate employment.

6. INDUSTRY RESULTS

The model in the preceding section affirms statistically the existence of the national jobless recovery. However, the model does not distinguish between competing causation theories. In this section, we address this issue by estimating industrially disaggregated versions of the model with structural breaks.²² The data used for this analysis consist of the monthly growth rates of employment in four sectors: durables manufacturing; nondurables manufacturing; trade, transportation, and utilities (TTU); and services.²³ Our focus is twofold. First, we are interested in determining which industries experienced statistically important changes in their business cycle dynamics. Second, we are interested in ascertaining whether the




FIGURE 3. Counterfactuals (payroll employment).

nature of these changes is consistent with the predictions of the aforementioned hypotheses about the jobless recoveries.

As before, Table 3 shows the Bayes factors for the break model against some alternatives for each dataset. Table 4 presents the summary statistics for the posterior distributions of the model parameters estimated with disaggregated data. Jobless recoveries are especially evident in the results for employment in both the durable and nondurable manufacturing industries. In this case, a single break estimated in December 1981 produces dramatic changes in the business cycle characteristics of both manufacturing employment series. As in the aggregate case, the speed of transition between regimes declines. However, the magnitude of this change is much larger (equating to approximately an 85% decline at the median) than for

TABLE 4. Posterior distributions for model parameters (industrially disaggregated)

Parameter	Model							
	Durables		Nondurables		TTU		Services	
	Pre	Post	Pre	Post	Pre	Post	Pre	Post
α_0	-1.84 (-2.81, -0.89)	-1.79 (-2.48, -1.12)	-1.34 (-2.09, -0.64)	-1.46 (-2.16, -0.88)	-0.47 (-1.10, -0.09)	-1.00 (-1.60, -0.50)	-0.64 (-1.46, -0.13)	-1.04 (-2.13, 0.05)
α_1	3.86 (2.57, 5.10)	3.15 (1.97, 4.45)	2.14 (1.19, 3.16)	2.59 (1.42, 3.98)	2.55 (1.77, 3.65)	2.65 (1.76, 3.74)	3.71 (2.54, 5.15)	3.46 (1.43, 5.53)
ϕ_{01}	0.35 (0.23, 0.47)	0.47 (0.31, 0.63)	0.58 (0.45, 0.72)	0.40 (0.22, 0.61)	0.22 (-0.09, 0.54)	0.42 (0.19, 0.66)	-0.20 (-0.98, 0.54)	0.41 (-0.80, 1.50)
ϕ_{11}	-0.31 (-0.48, -0.13)	-0.47 (-0.80, -0.17)	-0.41 (-0.60, -0.23)	-0.34 (-0.81, 0.06)	-0.09 (-0.53, 0.28)	-0.32 (-0.66, -0.01)	0.55 (-0.59, 1.81)	-0.50 (-2.70, 1.63)
ϕ_{02}	0.00 (-0.11, 0.11)	0.20 (0.04, 0.33)	-0.01 (-0.17, 0.15)	0.30 (0.14, 0.46)	0.18 (-0.15, 0.51)	0.15 (-0.03, 0.32)	0.33 (-0.41, 0.95)	0.19 (-0.88, 1.81)
ϕ_{12}	0.41 (0.22, 0.59)	0.44 (0.14, 0.79)	0.32 (0.09, 0.56)	0.08 (-0.26, 0.44)	0.13 (-0.27, 0.54)	0.26 (0.01, 0.54)	-0.01 (-0.93, 1.32)	0.34 (-2.28, 2.61)
σ^2	65.42 (57.82, 74.45)	10.72 (9.13, 12.67)	12.77 (11.24, 14.53)	3.76 (3.21, 4.52)	6.92 (6.12, 7.88)	3.81 (3.42, 4.27)	3.09 (2.68, 4.22)	2.57 (1.48, 3.00)
γ	2.27 (1.79, 2.32)	0.37 (0.19, 0.69)	2.40 (1.07, 3.36)	0.34 (0.19, 0.59)	0.85 (0.24, 2.83)	0.59 (0.34, 0.93)	0.15 (0.04, 0.44)	0.05 (0.01, 0.39)
c	-0.39 (-0.51, 0.04)	-0.46 (-0.90, -0.04)	-0.23 (-0.62, 0.18)	0.05 (-0.40, 0.47)	-0.04 (-0.46, 0.42)	0.02 (-0.39, 0.43)	0.05 (-0.31, 0.42)	0.04 (-0.46, 0.58)
d	3	3	2	3	4	3	3	3
τ	1981:12		1981:12		1981:12		1982:04	

Note: Intervals for 80% coverage in parentheses.

aggregate employment. Moreover, both the delay in response, d , and the threshold at which the response occurs, c , increase for nondurables. The top two panels of Figure 4 illustrate the implications of these changes for durables and nondurables, respectively. The decline in the rate of transition is manifested in two ways: (i) the recovery period is extended in the postbreak period and (ii) employment cycles are less pronounced, becoming closer to a pure autoregression.

For TTU, the estimated break date coincides with both manufacturing series. However, the estimated decline in the rate of transition is considerably smaller (approximately a 50% reduction at the mean and a 30% reduction at the median) than for either manufacturing sector. The resulting transition function shown in the third panel of Figure 4, then, retains its cyclical properties, albeit with longer recovery periods.

Finally, the services sector exhibits small quantitative and virtually no qualitative differences in its business cycle characteristics when estimated with a single endogenously chosen break. Although the median break date for services is slightly later (April 1982), each of the transition function parameters is statistically indistinguishable pre and postbreak. In fact, the rate of transition between expansion and recession for the services industry is sufficiently small so that the series exhibits virtually no business cycle characteristics. Indeed, the last panel of Figure 4 confirms that the transition function for services is nearly constant, resulting in a model quite similar to a linear AR. We therefore conclude that the jobless recovery appears to emanate from the manufacturing and—to some degree—trade sectors. Services appear not to be affected by whatever innovation spurred the change in the economy since the early 1980s.

We now examine these results in the context of the four hypotheses on the origin of the jobless recoveries. A consequence of the organizational restructuring hypothesis argued by Koenders and Rogerson (2005) is that jobless recoveries should be more prevalent in industries that experience longer expansions. Industries that expand longer should require more time to shed workers after a downturn. This would imply a correlation between the average length of an expansion period and the magnitude of the business cycle transition rate. A similar analysis can be made for the compositional changes in the labor supply hypothesis. If jobless recoveries are caused primarily by a shift to more self-employment after downturns, one would expect the change in transition dynamics to be statistically indistinguishable across industries. In other words, no industry should be affected more than another. Our results, however, do not support this hypothesis: We find, instead, that although each sector did experience longer expansions in the postbreak period, the business cycle dynamics of sectors were quite varied.²⁴

The most salient feature of industrially disaggregated models is the decline in cyclicity of the manufacturing industry. This result is consistent with theories in which jobless recoveries propagate from the manufacturing sectors, for example, just-in-time workers [Schreft and Singh (2003); Schreft et al. (2005)] or structural change [Groschen and Potter (2003)]. Distinguishing between these two hypotheses, however, requires an interpretation of the change in the manufacturing




FIGURE 4. Transition functions with industry-level employment. *Note:* Break date = December 1981 for 4a, 4b, 4c; Break date = April 1982 for 4d.

business cycle. In particular, the evidence here could be interpreted to suggest a move in the manufacturing sector from adjusting labor on extensive margins to adjusting on intensive margins. This interpretation is consistent with evidence for the automobile industry [Ramey and Vine (2006)], as well as other manufacturing sectors [Hetrick (2000)], and favors labor demand hypotheses in explaining the recent jobless recoveries.

7. CONCLUSIONS

In this paper, we explored the ability of empirical, nonlinear business cycle models to match the stylized facts of the recovery periods following the past two recessions. We have shown that a marriage between statistical models and jobless recoveries can be achieved by introducing a single endogenously chosen structural break into a smooth transition autoregression. After the break, the rate of transition between recession and expansion slows. For aggregate payroll employment, the transition rate falls approximately 60% equating to an extension of the average recovery period by at least one quarter. In addition, the resulting changes in business cycle dynamics are consistent with other evidence on the post-1980 period in the United States. Our results confirm a decline in both business cycle and idiosyncratic (high-frequency) volatility in the postbreak subperiod, consistent with the notion of a Great Moderation.

This decline in the responsiveness of employment to its lags is especially evident in some sectorally disaggregated employment series. Manufacturing, in particular, shows a dramatic decline in the rate of transition between recession and expansion after 1980, suggesting that the decline in cyclicity may have originated in the manufacturing industries. Lack of evidence for a change in transition dynamics in the services industry supports the hypothesis that changes in manufacturing practices (e.g., improved inventory management or adoption of information technologies) may have spilled over into manufacturing labor demand and provided the impetus for jobless recoveries. In particular, these innovations have disentangled the manufacturing industry's employment cycle from the output cycle, leading to the perception that employment growth now lags the output trough by an extended period.

NOTES

1. The strength of the evidence for jobless recoveries tends to vary depending on which employment survey is used. Studies contrasting payroll employment and the household survey suggest more concrete evidence from the former [Koenders (2005)].

2. The NBER turning points for the postwar period are November 1948–October 1949; July 1953–May 1954; August 1957–April 1958; April 1960–February 1961; December 1969–November 1970; November 1973–March 1975; January 1980–July 1980; July 1981–November 1982; July 1990–March 1991; and March 2001–November 2001.

3. Although the threshold model is not necessarily equivalent to the two-phase Markov-switching model employed by Hamilton (1989), the two exhibit similar behavior.

4. Because of the short recovery period between them, we combined the two recessions of the early 1980s for the figure.

5. In his overview of the latest jobless recovery, Schweitzer (2003) shows that payroll employment growth was slow. At the same time, the unemployment rate remained low and the labor force participation rate declined, which he partly attributes to the willing exit from the labor force of some participants from the 1990s expansion.

6. In a recent paper, Faberman (2008) notes that a “persistently high job destruction” rate followed the 1990–1991 recession and a “persistently low job creation” rate followed the 2001 recession. He relates the jobless recoveries to the Great Moderation, during which time the volatilities of those rates declined.

7. Aaronson et al. (2004b) argue that ignoring the difference between average growth and an industry’s growth could lead to an inaccurate assessment of sectoral reallocation. They examine various types of sectoral reallocation but do not find that it increased during the two jobless recoveries.

8. Similarly, Bachmann (2007) argues that, during the past two recovery periods, employers increased labor on the intensive margin (i.e., hours per worker) rather than the extensive margin (i.e., number of workers). He contends that this change in business practices accounts for about half of the differences in employment exiting the past two recessions.

9. Declining levels of business investment that have accompanied the persistently low levels of job creation might have played a role in the 2001 jobless recovery [Faberman (2004)].

10. Two papers attempt to model jobless recoveries explicitly. Glosser and Golden (2004) estimate a VAR and compute the impulse responses of employment to output shocks. Their model, however, has no explicit business cycle representation. Dueker (2006) estimates a four-regime Markov-switching model, in which one regime represents a high-growth bounce-back regime. He argues that this regime has been absent in the recovery periods of recent recessions. However, because the transition probabilities in the model are assumed to be independent, no explanation of why these periods are less prevalent is available.

11. A thorough review of the literature on Markov-switching models as they relate to business cycles can be found in Morley and Piger (2006). One paper has recently attempted to explain jobless recoveries in a Markov-switching framework. Holmes and Silverstone (2006) estimate the asymmetric responses of unemployment to changes in output across and within regimes. They find that unemployment responds more to a decrease in output than to an increase in output when the economy is in the above-trend unemployment regime, which might partially explain the two jobless recoveries.

12. For now, we neglect the possibility that the 1970 recession led to a jobless recovery [Koenders and Rogerson (2005)]. One could account for this possibility by estimating a model with Markov-switching transition function parameters or time-varying transition parameters. We leave this for future research.

13. We refer the reader to van Dijk et al. (2002) for a review of alternative representations for the transition function. In principle, we could test for the fit of the transition function [Escribano and Jorda (2001)]. However, the logistic function embeds some desirable features, including limiting equivalences discussed below.

14. For the estimation that follows, d_{\max} is taken to be the fourth monthly lag of employment growth.

15. We restrict the model with the monthly threshold variable to a single break between 1975:01 and 1990:12.

16. See Chib (1995) and Chib and Jeliazkov (2001) for a description of the algorithm used to compute the Bayes factors.

17. In all cases, the threshold variable is also annualized and seasonally adjusted.

18. For brevity, we forgo discussion of the household employment results and focus on results for payroll employment. We note cases in which there are qualitative differences in the results across surveys.

19. The median break dates are 1983:09 for payroll employment and 1983:11 for household employment.

20. In addition, we include the Bayes factors comparing the STAR model with and without a break to a benchmark linear autoregressive model. For aggregate employment data, the STAR model with a break is favored over the AR(2) model, but the AR(2) model is favored over the STAR model without a break.

21. The magnitude of the speed-of-adjustment parameter, γ , in all subsamples suggests that there exists some smooth transition dynamics.

22. We estimate each industry separately. An alternative specification might estimate the industries jointly in a smooth transition vector autoregression with an inverse-Wishart prior distribution on the covariance matrix. However, we forgo joint estimation in order to concentrate on the idiosyncratic variation in industry transition functions.

23. For each sector, the data sample ranges from 1963:01 to 2006:02. All data are seasonally adjusted and annualized. Each series is filtered to eliminate outliers larger than four standard deviations from the mean. These data are replaced with the average of six months of surrounding data.

24. A more thorough investigation of this hypothesis might include either more granularity or the addition of duration dependence in the transition function.

REFERENCES

- Aaronson, Daniel, Ellen R. Rissman, and Daniel G. Sullivan (2004a) Assessing the jobless recovery. *Federal Reserve Bank of Chicago Economic Perspectives* 28, 2–20.
- Aaronson, Daniel, Ellen R. Rissman, and Daniel G. Sullivan (2004b) Can sectoral reallocation explain the jobless recovery? *Federal Reserve Bank of Chicago Economic Perspectives* 28, 36–49.
- Bachmann, Ruediger (2007) Understanding Jobless Recoveries—A Tale of Two Margins. Manuscript, Yale University.
- Burns, Arthur F. and Wesley C. Mitchell (1946) *Measuring Business Cycles*. New York: National Bureau of Economic Research.
- Carlin, Bradley P., Alan E. Gelfand, and Adrian F. M. Smith (1992) Hierarchical Bayesian analysis of changepoint problems. *Applied Statistics* 41, 389–405.
- Chib, Siddhartha (1995) Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association* 90, 1313–1321.
- Chib, Siddhartha and Edward Greenberg (1995) Understanding the Metropolis–Hastings algorithm. *American Statistician* 49, 327–335.
- Chib, Siddhartha and Ivan Jeliazkov (2001) Marginal likelihood from the Metropolis–Hastings output. *Journal of the American Statistical Association* 96, 270–281.
- Dueker, Michael J. (2006) Using cyclical regimes of output growth to predict jobless recoveries. *Federal Reserve Bank of St. Louis Review* 88, 145–153.
- Escribano, Alvaro and Oscar Jorda (2001) Testing nonlinearity: Decision rules for selecting between logistic and exponential STAR models. *Spanish Economic Review* 3, 193–209.
- Faberman, R. Jason (2004) Gross Job Flows over the Past Two Business Cycles: Not All “Recoveries” Are Created Equal. Working Paper 372, U.S. Bureau of Labor Statistics.
- Faberman, R. Jason (2008) Job Flows, Jobless Recoveries, and the Great Moderation. Working Paper 08–11, Research Department, Federal Reserve Bank of Philadelphia.
- Gelfand, Alan E. and Adrian F. M. Smith (1990) Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* 85, 398–409.
- Glosser, Stuart and Lonnie Golden (2004) The changing nature of hours and employment adjustment in US manufacturing: A contributing cause of the jobless recovery? *International Journal of Manpower* 25, 618–642.
- Groschen, Erica L. and Simon Potter (2003) Has structural change contributed to a jobless recovery? *Federal Reserve Bank of New York Current Issues in Economics and Finance* 9, 1–7.
- Hamilton, James D. (1989) A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57, 357–384.

- Hetrick, Ron L. (2000) Analyzing the recent upward surge in overtime hours. *Monthly Labor Review* 123, 30–33.
- Holmes, Mark J. and Brian Silverstone (2006) Okun's law, asymmetries and jobless recoveries in the United States: A Markov-switching approach. *Economics Letters* 92, 293–299.
- Kim, Chang-Jin, James Morley, and Jeremy Piger (2005) Nonlinearity and the permanent effects of recessions. *Journal of Applied Econometrics* 20, 291–309.
- Koenders, Kathryn (2005) Long Expansions and Slow Recoveries: A Closer Look at Employment Fluctuations. Manuscript, Arizona State University.
- Koenders, Kathryn and Richard Rogerson (2005) Organizational dynamics over the business cycle: A view on jobless recoveries. *Federal Reserve Bank of St. Louis Review* 87, 555–579.
- Lopes, Hedibert F. and Esther Salazar (2006) Bayesian model uncertainty in smooth transition autoregressions. *Journal of Time Series Analysis* 27, 99–117.
- McConnell, Margaret M. and Gabriel Perez-Quiros (2000) Output fluctuations in the United States: What has changed since the early 1980's? *American Economic Review* 90, 1464–1476.
- Morley, James and Jeremy Piger (2006) The importance of nonlinearity in reproducing business cycle features. In Costas Milas, Philip Rothman, and Dick van Dijk (eds.), *Nonlinear Time Series Analysis of Business Cycles*, pp. 75–95. Amsterdam: Elsevier Science.
- Owyang, Michael T., Jeremy Piger, and Howard J. Wall (2008) A state-level analysis of the Great Moderation. *Regional Science and Urban Economics* 38, 578–589.
- Ramey, Valerie A. and Daniel J. Vine (2006) Declining volatility in the U.S. automobile industry. *American Economic Review* 96, 1876–1889.
- Schreft, Stacey L. and Aarti Singh (2003) A closer look at jobless recoveries. *Federal Reserve Bank of Kansas City Economic Review* 88, 45–73.
- Schreft, Stacey L., Aarti Singh, and Ashley Hodgson (2005) Jobless recoveries and the wait-and-see hypothesis. *Federal Reserve Bank of Kansas City Economic Review* 90, 81–99.
- Schweitzer, Mark (2003) Another jobless recovery? *Federal Reserve Bank of Cleveland Economic Commentary*, 1–4.
- Teräsvirta, Timo and Clive W. J. Granger (1993) *Modelling Nonlinear Dynamic Relationships*. New York: Oxford University Press.
- van Dijk, Dick, Timo Teräsvirta, and Philip Hans Franses (2002) Smooth transition autoregressive models—A survey of recent developments. *Econometric Reviews* 21, 1–47.