

Business Cycle Monitoring with Structural Changes*

Marcelle Chauvet and Simon Potter[†]

August 2008

Abstract

This paper examines the predictive content of coincident variables for monitoring U.S. recessions in the presence of instabilities. We propose several specifications of a probit model for classifying phases of the business cycle. We find strong evidence in favor of the ones that allow for the possibility that the economy has experienced recurrent breaks. The recession probabilities of these models provide a clearer classification of the business cycle into expansion and recession periods, and superior performance in the ability to correctly call recessions and to avoid false recession signals. Overall, the sensitivity, specificity, and accuracy of these models are far superior as well as their ability to timely signal recessions. The results indicate the importance of considering recurrent breaks for monitoring business cycles.

Keywords: Recession, Instability, Bayesian Methods, Probit model, Breaks.

*The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of New York or the Federal Reserve System.

[†]Address first author: Department of Economics, University of California, Riverside; email: chauvet@ucr.edu. Address second author: Macroeconomic and Monetary Studies, Federal Reserve Bank of New York, NY, 10045. Tel: (212) 720-6309; email: simon.potter@ny.frb.org.

1. Introduction

Since the seminal work of Burns and Mitchell (1946), a traditional method of monitoring the economy is through the use of coincident indicators. Burns and Mitchell classified hundred of macroeconomic variables into lagging, coincident, or leading according to the timing of their cyclical movements with the U.S. economic activity, in a research sponsored by the National Bureau of Economic Research (NBER). The same main coincident variables found in this study are still currently being used to date turning points of business cycles by the NBER Business Cycle Dating Committee, which is generally considered authoritative in dating recessions. These indicators are some of the most watched series by the press, businesses, policymakers, and stock market participants.

The reliability of predictions using these coincident variables in stationary models may have been compromised by the possibility of occurrence of structural breaks in the functioning of the economy in the last few decades. Several authors have found that the U.S. business cycle has experienced a substantial decline in its amplitude since the mid-1980s (e.g. McConnell and Perez-Quiros 2000, Koop and Potter 2000, Chauvet and Potter 2001, Van Dijk and Sensier 2004, among several others). An increased stability of business cycle fluctuations has important implications since it affects the frequency, duration, and probabilities of future recessions and expansions.

This paper examines the predictive content of coincident variables for monitoring U.S. recessions using several specifications of the probit model. First, we expand on previous research by considering probit specifications with fixed or endogenous breakpoints. In contrast with the research that investigates structural breaks in macroeconomic variables per se, we are interested in examining the stability of the relationship among the coincident variables and the business cycle. We find strong evidence of structural instability. Although the specification considered only a single endogenous break, there was substantial uncertainty about its location, spanning over five years. This is an important finding, given that the estimated probability of recessions is affected by both the presence as well as the location of the break.

Stabilization of the business cycle is not a new finding. The long expansion in the 1960s also spurred debates about economic stabilization. This is summarized in Arthur Burns' (1960) statement describing the stabilization of the U.S. economy since World War II:

“There is no parallel for this sequence of mild - or such a sequence of brief - contractions, at least during the past hundred years in our own country.”

Burns' statement refers to the decrease in volatility ('mildness') of the U.S. business cycle and to changes in the duration ('briefness') of business cycle phases after the War. A large literature followed studying these questions. The consensual evidence

was that an increased stabilization occurred in the U.S. economy comparing the periods before and after the War, but the magnitude of this stability was the subject of intense debates. Now, forty years later, economists are revisiting the evidence about the post-War business cycle stabilization - this time motivated by the finding of a structural break in the volatility of the U.S. output growth in the first quarter of 1984. Chauvet and Popli (2008) investigate whether this recent change is unique to the U.S. and particular to the 1980s or if it is part of a long run trend in volatility shared by several countries. They find strong evidence of multiple structural breaks leading to more stability in these countries over time, and that the recent decrease in U.S. output volatility is part of a broader long-term trend shared by all countries studied.

Based on these findings and on the large uncertainty regarding the location of the break in the probit specification studied, we propose an extension of the model that accounts for the potential existence of multiple shifts. We use business cycle troughs to date the breakpoints. Thus, assuming that business cycles are recurrent fluctuations, this is a model of recurrent breaks. We restrict the form of the break to be a change in the innovation variance implicit in the probit model. Most work using probit models in different applications assume stability and independent errors. We further extend the probit specification to allow for the possibility of serially correlated errors, introducing an autoregressive process for the latent variable. The models proposed are estimated using Bayesian methods and can be applied to different questions in a variety of fields that involve statistical surveillance, change point detection, tracking signals, among several others.

We find that the best fitting specification in terms of Bayes factor allows for recurrent breaks in the innovation variance as well as an autoregressive component. Analysis of the estimated probability of recessions shows that allowing for parameter changes increases substantially the signal to noise ratio. In particular, the recession probabilities for these models provide a clearer classification of the business cycle into expansion and recession periods and superior performance in the ability to correctly call recessions and to avoid false recession signals in-sample. Overall, the sensitivity, specificity, and accuracy (i.e., mean squared errors) of these models are far superior than for the models with a single break or with no break. On the other hand, the standard probit model with no break is the one with the lowest predictive ability. Finally, regarding the ability to timely signal recession, the model that considers multiple shifts in the innovation variances displays the best overall performance. These results indicate the importance of considering recurrent breaks for improving the monitoring of business cycles.

The issue studied is quite topical as there is currently a lot of speculation on whether the economy has recently entered into a recession. Most of the models considered indicate weaker economic activity in the end of the available sample of October 2007.¹

¹This paper was written in December 2007.

However, the value of the probability of a recession differs across specifications with weaker signals arising from the models with multiple shifts in the innovation variance.² On the other hand, the higher probability of recession from the probit model that considers no break or a single break reflects the uncertainty regarding the current state of the economy, and the importance of models with improved prediction ability. The prediction from the model with recurrent breaks is more in accord with the judgement that the economy was not in a recession in the end of 2007 by Martin Feldstein (2007), one of the members of the NBER dating Business Cycle Dating Committee.³

The paper is organized as follows. Section two introduces the probit model and discusses our various extensions. The third section describes the Gibbs sampler. In the fourth section the empirical results from the alternative specifications are presented. The fifth section concludes.

2. The Models

2.1. Standard Probit Model

The standard probit model assumes an unobservable variable Y_t^* with a corresponding binary indicator Y_t for whether Y_t^* is positive or negative:

$$Y_t = \begin{cases} 0 & \text{if } Y_t^* < 0 \\ 1 & \text{if } Y_t^* \geq 0 \end{cases} . \quad (2.1)$$

In our context, Y_t^* represents the state of the economy as measured by the NBER recession dating: Y_t takes the value 0 if the observation is an expansion or 1 if it is a recession. The latent variable Y_t^* is related to the regressors X_{it} according to the equation:

$$Y_t^* = \beta_0 + \beta' X_t + \varepsilon_t \quad \varepsilon_t | X_t \sim i.i.d.N(0, 1) \quad (2.2)$$

²The models are devised to monitor current economic conditions as they only use coincident (not leading) variables. Thus, the models do not have information on future probabilities of a recession from November 2007 on.

³Martin Feldstein (December 16, 2007):

”Because monthly data for December won’t be available until next year, we cannot be sure whether the economy has turned down. The measure of personal income for October suggests that the economy may have peaked and begun to decline, but the data for employment and industrial production in November and for sales in October show continued growth. My judgment is that when we look back at December with the data released in 2008 we will conclude that the economy is not in recession now.”

where $X_t = \{X_{it} : i = \text{Prod, Sales, Inc, and Emp}\}$ are the coincident macroeconomic variables industrial production, sales, personal income, and employment, and $\beta = \{\beta_i : i = \text{Prod, Sales, Inc, and Emp}\}$ are regression coefficients. The model assumes that:

$$P(Y_t^* \geq 0 | X_t, \beta) = \Phi[\beta_0 + \beta X_t], \quad (2.3)$$

where Φ is the cumulative distribution function of the standard normal distribution, $\beta = [\beta_0, \beta]$, and $P(Y_t^* \geq 0 | X_t, \beta)$ is the conditional probability of a recession. This standard model can be interpreted as capturing the behavior of the business cycle dating committee in an economy with no changes.

2.2. Extensions of the Probit Model

We extend the standard probit model in three ways. First, we allow for the presence of a fixed or endogenous breakpoint. Second, we allow the variance of the innovation to change with the business cycle, which accounts for recurrent breaks in the innovation variance. Finally, we allow for the possibility of serially correlated error by adding an autoregressive component to the model. These alternative specifications imply very different predictive ability, as discussed in section 4. All models are estimated using Bayesian methods.

2.2.1. Probit Model with a Breakpoint

In order to allow for structural breaks, the unobserved Y_t^* is modeled as a Gaussian process with constant unit innovation variance and conditional mean:

$$E(Y_t^*) = \beta(\tau)' \mathbf{Z}_t(\tau), \quad (2.4)$$

where τ represents a break point and $\mathbf{Z}_t(\tau)$ is the 10×1 vector:

$$\mathbf{Z}_t(\tau) = [1(\tau_1), 1(\tau_2), 1(\tau_1)X_{it}, 1(\tau_2)X_{it}]',$$

For $\tau_1 \in \{t : 1 \leq t \leq \tau\}$ and for $\tau_2 \in \{t : \tau < t \leq T\}$. Let the collection of parameter vectors $\{\beta(\tau) \in [\tau_1, \tau_2]\}$ be defined by:

$$[\beta_0(\tau_1), \beta_0(\tau_2), \beta(\tau_1), \beta(\tau_2)]'.$$

The date τ can be fixed or be endogenously estimated as explained in the next section. The use of Bayesian methods allows us to capture the joint uncertainty over the timing of the breakpoint and the parameter estimates.

2.2.2. Probit Model with Recurrent Breakpoints

If Y_t^* is multiplied by any positive constant, the indicator variable Y_t is not changed, which implies that the coefficients $\{\beta_0, \beta\}$ can be estimated only up to a positive multiple. Thus, the standard probit model assumes that the variance of the errors is equal to one to fix the scale of Y_t^* . We consider a more general specification in which the variance of the innovation may change:

$$Y_t^* = \beta_0 + \beta' X_t + \sigma(t)\varepsilon_t, \quad (2.5)$$

and the initial business cycle is partially observed starting at $t = 1$ ($t_1 = 0$). We restrict this specification by assuming that the innovation variance is the same within each business cycle.

We assume that a business cycle starts the month after a NBER trough and continues up to the month of the next NBER trough. If t_{n-1} corresponds to the beginning of business cycle $n - 1$, then the dates of business cycle n are $t \in \{t_{n-1} + 1, \dots, t_n - 1, t_n\}$. Let the periods in which the economy is in an expansion or recession of business cycle n be denoted by the sets E_n and R_n , respectively. Hence, business cycle expansions and recessions are classified by $E = \cup E_n$ and $R = \cup R_n$, respectively, where

$$\sigma_n = \sigma(t) \text{ if } t_{n-1} < t \leq t_n, n = 1, \dots, N,$$

Since the scale of the innovation and the coefficient parameters cannot be separately identified, this can be interpreted as a restricted time-varying parameter specification, in which the innovation variance is normalized to 1 across all business cycles, but each cycle has a unique intercept $\beta_{n0} = \beta_0/\sigma_n$ and slope coefficients $\beta_{ni} = \beta_i/\sigma_n$. Another interpretation of the model is that the scale of shocks themselves may change across business cycles. This allows for the possibility that the size of the innovation variance may change depending on the duration of the business cycle. More importantly, since it allows for potential recurrent breakpoints across business cycles, it can capture long run trends on the variance such as whether it has been decreasing over time. For example, suppose that the change in the scale of the shocks is common across the business cycle variables in X_t and is the same as the change in scale to the innovation to the latent variable. In this case our approach could alternatively be interpreted as keeping constant the scale of the business cycle variables to maintain the same relationship to business cycle phases. Thus, in this model we assume that:

$$P(Y_t^* \geq 0 | X_t, \boldsymbol{\beta}) = \Phi_n[(\beta_0 + \beta' X_t)] = \Phi[(\beta_0 + \beta' X_t)/\sigma_n]. \quad (2.6)$$

We assume in the estimation below that σ_n are a priori independent across business cycles, as in Chauvet and Potter (2005).⁴

⁴An interesting extension is proposed in Koop and Potter (2007), who allow for dependence across

2.2.3. Probit Model with Multiple Breakpoints and Dependence in the Latent Variable

We further extend the probit model to allow for the possibility that the latent variable Y_t^* follows a first order autoregressive process:

$$Y_t^* = \beta_0 + \beta' X_t + \theta Y_{t-1}^* + \sigma(t)\varepsilon_t, \quad (2.7)$$

where the autoregressive parameter $|\theta| < 1$. The intention here is to capture dependence in the business cycle phases that comes from the concept that recessions and expansions should have a certain duration. In our previous work on recession forecasting we have found that allowing for dependence can eliminate a number of false positives.

2.2.4. Likelihood Function

The likelihood function for the most general model that consider recurrent breakpoints in the variances, a single breakpoint in the conditional mean, and autoregressive process for the latent variable is:

$$\begin{aligned} \ell(Y^T | X_i^T, \beta, \{\sigma_n\}, \theta, y_0) \\ = \prod_{n=1}^N \left\{ \prod_{t \in R_n} \Phi_n[\beta(\tau)' \mathbf{Z}_t(\tau) + \theta Y_{t-1}^*] \prod_{t \in E_n} (1 - \Phi_n[\beta(\tau)' \mathbf{Z}_t(\tau) + \theta Y_{t-1}^*]) \right\}. \end{aligned} \quad (2.8)$$

All of the other models considered are nested within this likelihood function with appropriate parameter restrictions.

2.2.5. Models Estimated

We estimate five different specifications:⁵

1. Model 1: Standard probit with no break;
2. Model 2: Probit model with a fixed break in 1984:01;
3. Model 3: Probit model with an endogenous break;
4. Model 4: Probit model with recurrent breaks - business cycle specific variance;
5. Model 5: Probit model with recurrent breaks and autoregressive latent process.

business cycles. One could also consider the unrestricted model with time variation each month, as suggested by the editor. We are implementing this in an on-going extension of our framework by allowing for stochastic volatility in the innovation variance with the initial value always set to unity.

⁵We opted not to estimate an model with an endogenous break in the conditional mean and multiple breaks with an autoregressive latent process because of the computational difficulties we encountered with Model 3.

3. Estimation Method

We use Bayesian methods to evaluate the posterior properties of the probit specifications. The Bayesian techniques used have several advantages: the Gibbs sampler can be used to simulate the latent variable, which simplifies considerably the computation of the likelihood function and of Bayes factors used to compare the various models. This can be implemented using the Savage-Dickey Density ratio within the Gibbs sampler; the distributions over the probability predictions contain information on uncertainty regarding parameters, existence or location of breakpoints, and over the most recent value of the latent variable. Latent variables determine the existence and location of the break. Conditional on the latent variable the model parameters are generated, while conditional on the parameters, the latent variable is generated. These simulations are repeated until an adequate sample is obtained from the posterior distributions.

3.1. Models with a Single Breakpoint

For model 3, we assume that the prior on $\beta(\tau)$ is identical and independent across different breakpoints. Since we need a non-diffuse prior in order to calculate the marginal likelihood, we assume that the prior distributions of $\beta(\tau)$ are normal with mean $\{\underline{\beta}\}$ and variance $\underline{\mathbf{D}}$. The prior distribution π on τ is assumed to be discrete uniform over support $[t_1, t_2]$.

The Gibbs sampler procedure generates draws of the latent Y_t^* conditional on a draw of $\{\beta(\tau), \tau\}$. The sampler is:

1. Draw $Y_t^* < 0$ by adding a draw from the truncated normal on $(-\infty, -\mathbf{Z}_t(\tau)' \beta(\tau))$ to $\mathbf{Z}_t(\tau)' \beta(\tau)$ if t is an expansion period.
2. Draw $Y_t^* \geq 0$ by adding a draw from the truncated normal on $(-\mathbf{Z}_t(\tau)' \beta(\tau), \infty)$ to $\mathbf{Z}_t(\tau)' \beta(\tau)$ if t is a recession period.

Given this sequence of draws for Y_t^* we can then construct the marginal likelihood for each possible breakpoint. Defining $\mathbf{Y}^* = [Y_1^*, \dots, Y_T^*]'$ and $\mathbf{Z}(\tau) = [Z_1(\tau)', \dots, Z_T(\tau)']$ we have:

$$f(\mathbf{Y}^* | \tau) = \left[\frac{\det(\overline{\mathbf{D}}^{-1}(\tau))}{\det(\underline{\mathbf{D}}^{-1})} \right]^{-0.5} \times \exp \left[\begin{array}{c} -0.5(vs^2 + (\overline{\beta}(\tau) - \widehat{\beta}(\tau))\mathbf{Z}(\tau)'\mathbf{Z}(\tau)(\overline{\beta}(\tau) - \widehat{\beta}(\tau))' \\ + (\overline{\beta}(\tau) - \underline{\beta})\underline{\mathbf{D}}^{-1}(\overline{\beta}(\tau) - \underline{\beta})' \end{array} \right],$$

where $vs^2 = (\mathbf{Y}^* - \mathbf{Z}(\tau)\widehat{\beta}(\tau))'(\mathbf{Y}^* - \mathbf{Z}(\tau)\widehat{\beta}(\tau))$ is the sum of the squared errors, $\overline{\mathbf{D}}(\tau) = [\underline{\mathbf{D}}^{-1} + \mathbf{Z}(\tau)'\mathbf{Z}(\tau)]^{-1}$ is the posterior variance, $\overline{\beta}(\tau) = \overline{\mathbf{D}}(\tau) [\underline{\mathbf{D}}^{-1}\underline{\beta} + \mathbf{Z}(\tau)'\mathbf{Y}^*]$

is the posterior mean, and $\widehat{\boldsymbol{\beta}}(\tau) = [\mathbf{Z}(\tau)' \mathbf{Z}(\tau)]^{-1} \mathbf{Z}(\tau)' \mathbf{Y}^*$ is the ordinary least squares estimator.

We find the marginal likelihood for each potential breakpoint τ , and then draw a new value of τ from the probability mass function:

$$\frac{f(\tau | \mathbf{Y}^*)}{\sum_{\tau=t_1}^{t_2} f(\tau | \mathbf{Y}^*)}.$$

Given the draw of τ , we draw $\boldsymbol{\beta}(\tau)$ from the normal distribution with mean $\overline{\boldsymbol{\beta}}(\tau)$ and variance $\overline{\mathbf{D}}(\tau)$.

3.2. Models with Recurrent Breakpoints

The model with an autoregressive process requires multiple integration over the unobserved lagged variable. We use the Gibbs sampler to evaluate the likelihood function as described below.⁶

3.2.1. Obtaining Draws of the Latent Variable

The Gibbs sampler starts by generating draws of the latent Y_t^* conditional on $(\beta_0, \beta, \{\sigma_n\}, \theta)$ and the observed coincident variables. Let $X_t' \boldsymbol{\beta} = \beta_0 + \beta' X_t$. If $\theta = 0$, then the sampler would have the following simple form:

1. Draw ε_t from the truncated normal on $(-\infty, -X_t' \boldsymbol{\beta} / \sigma_n)$ if t is an expansion period of business cycle n .
2. Draw ε_t from the truncated normal on $[-X_t' \boldsymbol{\beta} / \sigma_n, \infty)$ if t is a recession period of business cycle n .

The lagged value of the latent variable in the conditional mean makes the sampler more complex. Consider first generating the last value in the observed sample, Y_T^* . If we could condition on a value for Y_{T-1}^* , then we could use the steps above by redefining $X_T' \boldsymbol{\beta} = \beta_0 + \beta' X_{iT} + \theta Y_{T-1}^*$. This would generate a draw of the last period value of the latent variable. With this ‘new’ value of Y_T^* and the ‘old’ value of Y_{T-2}^* , we can use the *a priori* joint normality of the underlying latent variable model to form a conditional normal distribution for Y_{T-1}^* . The exact form of this distribution depends on an assumption about the initial value Y_0^* . We simplify the analysis by assuming that

⁶Geweke (1999) and Chib (2001) are seminal work on modern Bayesian computational techniques. These papers, their reference to earlier work, as well as Dueker (1999, 2001) are closely related to the techniques we use to estimate the models here and in Chauvet and Potter (2002, 2005).

$Y_0^* = \beta_0 + \beta' X_{i0} = 0$. Then, as shown in the appendix, we have *a priori* the conditional normal distribution with mean:⁷

$$\tilde{X}_{T-1} + \theta \begin{bmatrix} V(\tilde{Y}_{T-1}) \\ V(\tilde{Y}_{T-2}) \end{bmatrix} \begin{bmatrix} V(\tilde{Y}_T) & \theta^2 V(\tilde{Y}_{T-2}) \\ \theta^2 V(\tilde{Y}_{T-2}) & V(\tilde{Y}_{T-2}) \end{bmatrix}^{-1} \begin{bmatrix} Y_T^* - \tilde{X}_T \\ Y_{T-2}^* - \tilde{X}_{T-2} \end{bmatrix},$$

and variance:

$$V(\tilde{Y}_{T-1}) - \theta^2 \begin{bmatrix} V(\tilde{Y}_{T-1}) \\ V(\tilde{Y}_{T-2}) \end{bmatrix} \begin{bmatrix} V(\tilde{Y}_T) & \theta^2 V(\tilde{Y}_{T-2}) \\ \theta^2 V(\tilde{Y}_{T-2}) & V(\tilde{Y}_{T-2}) \end{bmatrix}^{-1} \begin{bmatrix} V(\tilde{Y}_{T-1}) \\ V(\tilde{Y}_{T-2}) \end{bmatrix}',$$

where

$$V(\tilde{Y}_t) = \sum_{s=0}^{t-1} \theta^{2s} \sigma^2(t-s)$$

and

$$\tilde{X}_t = \sum_{s=0}^{t-1} \theta^s X'_{t-s} \beta.$$

Hence, we draw from the appropriate truncated normal above to obtain a new draw of Y_{T-1}^* . This procedure is repeated until we get to the initial observation period. The value of Y_1^* is drawn in a similar manner to Y_T^* , by conditioning on the new draw of Y_2^* . However, this value has a different form, since its mean is given by:

$$\tilde{X}'_1 \beta + \frac{\theta \sigma^2(1)}{\sigma^2(1) + \theta^2 \sigma^2(2)} (Y_2^* - \tilde{X}'_2),$$

and variance by:

$$\sigma^2(1) - \frac{\theta^2 \sigma^4(1)}{\sigma^2(1) + \theta^2 \sigma^2(2)}.$$

3.2.2. Obtaining Draws of the Model Parameters

Once obtained the sequence of draws for $\{Y_t^*\}$, we can then generate draws of the parameters $(\beta, \{\sigma_n\}, \theta)$. We use as priors the normal and gamma distributions, which generate simple conditional distributions for the posterior. We assume that the parameters of the conditional mean β are *a priori* bivariate normal with mean vector $\mu(\beta)$ and variance matrix $V(\beta)$. In addition, the autoregressive parameter θ is assumed to have an *a priori* truncated normal on $(-1, 1)$ with mean $\mu(\theta)$ and variance $V(\theta)$ independent

⁷Dueker (1999) derives this conditional distribution using an alternative representation. Our choice of this representation is based on the fact that it shows directly the dependence on the history of the coincident variables and the business cycle specific variances.

of β . Finally the $N - 1$ variance parameters are assumed to be *a priori* independent with identical inverted gamma distributions.

For a draw of β , define the time series $W_t^\beta = (Y_t^* - \theta Y_{t-1}^*)$. Then, conditional on $(\{Y_t^*\}, \{\sigma_n\}, \theta)$, the parameters β are obtained from a normal distribution with variance matrix:

$$V(\bar{\beta}) = \left[V(\underline{\beta})^{-1} + \sum_{t=1}^T X_t X_t' / \sigma^2(t) \right]^{-1},$$

and mean vector:

$$\mu(\bar{\beta}) = V(\bar{\beta}) \left[V(\underline{\beta})^{-1} \mu(\underline{\beta}) + \sum_{t=1}^T X_t W_t^\beta / \sigma^2(t) \right].$$

For a draw of θ , define the time series $W_t^\theta = (Y_t^* - X_t' \beta)$. Then, conditional on $(\{Y_t^*\}, \{\sigma_n\}, \beta)$ a potential draw for θ is from a normal distribution with variance:

$$V(\bar{\theta}) = \left[V(\underline{\theta})^{-1} + \sum_{t=1}^T Y_{t-1}^{*2} / \sigma^2(t) \right]^{-1},$$

and mean:

$$\mu(\bar{\theta}) = V(\bar{\theta}) \left[V(\underline{\theta})^{-1} \mu(\underline{\theta}) + \sum_{t=1}^T Y_{t-1}^* W_t^\beta / \sigma^2(t) \right].$$

We keep generating draws until the stationary condition is satisfied.

For draws of $\{\sigma_n\}$, we assume that the prior distributions are independent inverted gammas with identical degrees of freedom $\underline{\nu}$ and scale $\underline{\nu s}^2$. Hence, the prior mean is $\underline{\nu s}^2 / (\underline{\nu} - 2)$. The prior parameters are then updated for the business cycle $n \geq 2$ by:

$$\begin{aligned} \bar{\nu}_n &= \underline{\nu} + t_n - t_{n-1} \\ \bar{\nu s}_n^2 &= \underline{\nu s}^2 + \sum_{t=t_{n-1}+1}^{t_n} (Y_t^* - X_t' \beta - \theta Y_{t-1}^*)^2. \end{aligned}$$

3.3. Evaluating Different Models

As discussed above we can calculate Bayes factor for the different models by the Savage Dickey Density Ratio since the likelihood functions are nested (see Koop and Potter 1999).⁸ The Bayes factor gives information on how well each model fits over the observed sample. Its advantages over the likelihood ratio lie on how it deals with nuisance

⁸An earlier version of the paper contained details on the calculations of Bayes Factors for these models.

parameters and averages over possible parameter values. These advantages are especially important in assessing whether to choose a model with breaks over one without, although for the purposes of monitoring a particular phase of the business cycle it is perhaps more important to consider measures of predictive performance. This latter point is particularly relevant for model 2, for which the evidence collected by the Bayes factor before 1984 is irrelevant to any current prediction or monitoring problem.

4. Empirical Results

We consider five specifications of the probit model as described in section 2.2.5. First, we estimate the probit models under the assumption of no structural breaks using both maximum likelihood methods and the Gibbs sampler. We then start the sampler from the maximum likelihood estimator. For the Gibbs sampler we use 100,000 iterations, but calculate the posterior properties only after 10,000 draws (i.e., we use 90,000 iterations to estimate the probit models). Next, we fix a breakpoint in the conditional mean of the probit model in January 1984, based on the evidence of increased stability in the U.S. economy from this date onwards, as documented by several authors. Third, we consider estimating the breakpoint endogenously. Fourth, we allow for the possibility of recurrent breaks in the variance of the model across business cycles. Finally, we consider a version of the probit model with changing variance and autocorrelated latent variable.⁹

4.1. Priors

Although the influence of the prior is very minor given the sample size, we give to the standard probit model 1 the advantage of a prior centered at its maximum likelihood estimator. For all models we assume that the prior variance of β is the identity matrix ($\beta(\tau)$ for models 2 and 3) and the prior mean is the maximum likelihood estimate from the benchmark model 1, which assumes unit variance, no breaks, and no autoregressive component. This choice has no effect on the results (as we have almost 600 observations). In addition, it simplifies the calculation of the Bayes factors and is preferable to a prior that centered β at zero, which would imply no knowledge that we are using business cycle variables. It also implicitly favors the simplest model as its prior is aligned with the maximum likelihood estimates. With respect to the prior for τ , all possible datapoints could be used since we are assuming informative priors. However, we decided to keep a minimum of 10% of the data in each regime, eliminating a proportion of the endpoints. Notice that this sample still allows for the possibility of a break after the last business cycle trough of November 2001. It is important to note that by imposing a

⁹The computation time for the endogenous break model 3 is around 10 hours, whereas it takes only 2 hours for the most complicated model 5.

breakpoint in Model 2 we are effectively estimating it using two different samples. This provides a benchmark for the endogenous approach proposed. We could have made as the breakpoint the start of the 1980s expansion, but we decided to use the standard finding in the literature of a break in early 1984. Finally, the prior for the autoregressive parameter θ is assumed to be a truncated standard normal on $(-1, 1)$, and the prior for the business cycle specific variances in models 4 and 5 is centered at 1 with diffuse degrees of freedom equal to 3.

4.2. Data

Based on the seminal work of Burns and Mitchell (1946), the NBER Business Cycle Dating Committee considers four main monthly indicators in determining business cycle chronology: industrial production (Production), real manufacturing and trade sales (Sales), real personal income less transfer payments (Income), and employees on non-agricultural payroll (Employment). We use the same coincident variables as the NBER except for employment. We follow Chauvet (1998) and Chauvet and Hamilton (2005) in using instead civilian labor force in nonagricultural industries (TCE) rather than employees on nonagricultural payrolls (ENAP) as used by the NBER. ENAP is based on a survey of business establishments, whereas TCE is based on a survey among households. These two employment series have generally moved together, with some minor differences around business cycle turning points. Although the revised ENAP may reflect better labor conditions ex-post, its performance in capturing real time cyclical changes in the economy is weaker compared to the household survey (TCE). In fact, ENAP tends to lag business cycle in real time, whereas TCE coincides with business cycle phases and calls turning points a lot faster, which will be an important feature in the analysis of the probability of a recession in our endpoint.

All series are transformed as 100 times their log first difference and the sample available is from February 1959 to October 2007.

4.3. Results

Table 1 reports the posterior means of the coefficients for the five models considered. The results indicate a significant relationship between a decrease in the coincident variables and the probability of a recession.

The coefficients in models 2 and 3 are substantially different before and after break. In particular, the post-break parameters are a lot higher in absolute values. Recall that since the scale of the innovation and the coefficient parameters β are not separately identified (i.e, $\beta^{\tau_0} = \beta/\sigma_{\tau_0}$ and $\beta^{\tau_1} = \beta/\sigma_{\tau_1}$), this could possibly be reflecting the fact that the innovation variance has decreased since the break. The coefficients from model 1 are roughly an average of the coefficients for the periods pre and post break in models

2 and 3. Nevertheless, the results from models 1, 2, and 3 differ in several ways.

First, we find strong evidence of the existence of a structural break in the relationship between the monthly coincident series and the business cycle with the natural log of the Bayes factor equals to -295.5 (model 3 vs model 1). However, even with 90,000 iterations the Bayes factor is not precisely estimated, which suggests uncertainty regarding the location of the breakpoint. Figure 1 plots the posterior break probability obtained from model 3, which endogenously estimates the breakpoint. In contrast to previous studies that find a break in 1984 in the volatility of GDP and other macroeconomic series, we find that the change in the *relationship* between the coincident series and the business cycle phase indicator occurred between 1977 and 1982. Thus, fixing a breakpoint in 1984 as in model 2 is not appropriate for monitoring business cycles with the series studied, even though we know it captures the decline in variation in many business cycle variables including the covariates used in our study. In addition, the spanning of the posterior distribution of the probabilities of a break over five years indicates considerable uncertainty about its exact location.

The posterior means of the conditional mean parameters from models 4 and 5 can not be directly compared to the others since they assume changing variance across business cycles and, in the case of model 5, an autoregressive component is also included. Table 1 shows the different values obtained for the variance across the six complete business cycles in the sample for these models. The estimated variances reflect similar findings across different business cycles, with the highest value occurring in the shortest business cycle in 1980-1982, which is four times the average variance. This corresponds to the period in which the Federal Reserve changed its operating procedures between 1979-1982, and the US economic experienced high volatility. Notice that this is also the period for which model 3 indicates a structural break. However, the Bayes factor strongly favors recurrent breaks (model 4) over a single endogenous break (model 3), with $\ln BF = -29.4$. Using Jeffrey's (1961) rule, the factors indicate a decisive evidence against the null.¹⁰

On the other hand, the variances with lowest values are associated with the long expansions of the 1960s, 1980s, 1990s, and 2000s. An interesting feature shown in the business cycle specific variance models is that the innovation variance shows a declining trend over time, especially since 1982. In fact, apart from the volatile time during the 1975-1980 business cycle, the variance has been decreasing since the beginning of the sample. Thus, imposing constant variance across business cycles can be misleading.

The consideration of recurrent shifts is important as the parameters change considerably as shown in Table 1. In addition, where the break occurs is important in

¹⁰ Jeffrey's rule assesses evidence against the null as follows: $\ln BF > 0$ evidence supports null; $-1.15 < \ln BF < 0$ very slight evidence against null; $-2.3 < \ln BF < -1.15$ slight evidence against the null; $-4.6 < \ln BF < -2.3$ strong to very strong evidence against the null; $\ln BF < -4.6$ decisive evidence against null.

determining the probabilities of recession. In particular, the uncertainty over the breakpoints implies very different signal to noise ratio regarding prediction of recessions. The posterior mean of the probabilities of recession for the probit models is given by $\Phi(\cdot)$ for draws of the parameters, and is plotted in Figure 2 (full sample) and Figure 3 (last two decades). The probabilities consistently rise before each of the seven recessions in the sample as dated by the NBER. However, there are marked differences among the signal to noise ratios of each model. Although the probabilities are somewhat similar in the first part of the sample across models 1, 2, and 3, the specifications assuming a known or endogenous breakpoint obtain different predictions. Model 1 is noisy during expansions and signals recessions with relatively low probability values. On the other hand, the performance of model 3 with endogenous break is better than the simpler versions. This is also the case for models 4 and 5 compared to the other models, which show a much clearer dichotomy between recessions and expansions and, therefore, less uncertainty regarding interpretation of these probabilities.

This evidence is formalized in Table 3, which reports the predictive ability of the models in correctly signaling recessions and expansions with a cutoff of 50% (classification table). It is important to note that this is not a pseudo real time exercise since we use full sample estimates and, perhaps more importantly, we do not use the vintage of the data available to analysts during these periods. With these caveats it can be seen that all models show good performance in predicting expansions with specificity (percentage of correct expansion predictions) of 97% and above. The best performance is obtained for model 4 with a rate of 99.6%, followed by model 5, with 99.4%. Model 5 has an additional advantage since the autoregressive latent variable is found using the whole sample information. Nevertheless, its ex-post classification ability is very impressive.

Expansion phases are long in the U.S. while recessions are short and abrupt events, with a minimum duration of 6 months¹¹ and a maximum length of 16 months, which occurred in 1982. On the other hand, expansions have a minimum duration of one year and an average length of 57 months in the sample studied. The duration of these phases makes it easier to correctly predict expansions than recessions. Consequently, the ability to predict recessions is more variable across specifications. Model 1 exhibits the worst sensitivity performance (percentage of correct recession predictions), with a rate of 29%. The consideration of a breakpoint increases substantially the sensitivity of models 2 and 3 compared to model 1, with a rate of 40% and 45%, respectively. Note that endogenously estimating the breakpoint improves the ability of the model to predict recessions compared to imposing a fixed breakdate in 1984.

Finally, the models with recurrent breaks across business cycles are the ones that

¹¹The NBER defines recession as a broad contraction of the economy with a minimum duration of 6 months.

present the best performance in predicting recessions. The sensitivity of model 4 is 71%, and of model 5 is 84%. We also compare the predictive ability of the specifications with a naive model that assumes constant probability. The total percentage gain of model 1 over the naive model is only 18%, while it is 22% for model 2, and 30% for model 3. Once again, models 4 and 5 display far superior performance than the other models with an overall percentage gain over the naive model of 68% and 80%, respectively. These numbers are impressive and indicate the importance of considering recurrent breaks for improving the predictive performance in the probit models.

From the evidence above, the specification that considers both business cycle-specific variance and autoregressive parameter (model 5) is the one with better sensitivity and percentage gain. Figures 2 and 3 show that the posterior mean probabilities of recession for this model are very smooth, with very low noise during expansions. Note that although the probabilities consistently increase above 70% *during* each recession in the sample, the model tends to oversmooth the signals at the beginning of recessions (i.e. peaks) yielding delayed recession calls. Table 3 shows the lead and lag signals of recessions for the alternative specifications. Model 5 is the one with worst performance, consistently calling recessions with delays longer than models 2, 3, and 4. Under this criterion, model 4 is the one with best overall ability to timely signal recessions. This is also confirmed by Yates' (1982) decomposition in Table 4, which shows that model 4 has the highest accuracy rate (lowest mean squared error).

4.3.1. Current Recession Probabilities

The last observation available as of December 2007 (when this article was originally written) is for October 2007. Over the early fall of 2007 the perception that the US might have been entering a recession increased considerably. The probability of recession for October 2007 is 39% for model 1 and 37% for model 3. This probability is even higher for model 2, 54%, which indicates the beginning of a recession under a cutoff of 0.5. However, model 4 indicates a much smaller probability for this month, 22% and for model 5 this probability is only 1%. Given the different performance of the alternative models in correctly and timely signalling a recession, the posterior probabilities do not give much information on the uncertainty regarding these values. Figure 4 shows the posterior cumulative distribution function of the probability of a recession state in October 2007. Under the assumption of a break in 1984 (model 2), 95% of the posterior on the probability of recession is between 0.3 and 0.7, while if one assumes recurrent breaks (model 4), 95% of the posterior on the probability of a recession is between 0.02 and 0.28. That is, the uncertainty regarding the recession probability decreases substantially when taking into account business cycle-specific variances in model 4. For model 5, 95% of the posterior on the probability of a recession is between 0 and 0.01, but given its delay in calling recessions, the results from model 4 are more reliable based

on previous performance.¹²

Note that the models examined do not have information on the probability of a recession in November 2007 and subsequent months, since they are devised to monitor current economic conditions as they only use coincident (not leading) variables. This on its own carries a lot of uncertainty and it is an important task since the NBER takes between 6 months to 20 months to announce the beginning or end of a recession after the fact.

5. Conclusions

This paper extends a standard probit specification for monitoring business cycles to account for the possibility of single breakpoint or recurrent shifts and serially correlated errors. A Gibbs sampling algorithm is used for estimating the effects of breaks on the estimated probability of recession.

We find strong evidence of the existence of a break in the relationship between the monthly coincident series and the business cycle. However, the results suggest considerable uncertainty about its exact location and gives support to the assumption of recurrent breaks. This is confirmed by the Bayes factor and the superior classification performance of models that allow for recurrent shifts in the innovation variance. The recession probabilities for these models provide a clearer classification of the business cycle into expansion and recession periods, and superior performance in the ability to correctly call recessions and to avoid false recession signals in-sample. The results indicate the importance of considering recurrent breaks for monitoring business cycles. In an on-going research we are studying ways to use these models for real-time classification with limited information on the dating decisions of the NBER business cycle committee.

Appendix

The formula required for the Gibbs sampler draws for models 4 and 5 is derived below. The full conditional distribution under the first order autoregressive assumption

$$f(Y_t^* | Y_T^*, \dots, Y_{t+1}^*, Y_{t-1}^*, \dots, Y_{K+1}^*)$$

is equivalent to:

$$f(Y_t^* | Y_{t+1}^*, Y_{t-1}^*).$$

Since $Y_{t+1}^*, Y_t^*, Y_{t-1}^*$ have a joint normal distribution, the conditional distribution is normal. Under the assumption that all initial values are zero, we can write the latent

¹²The low probabilities of recession from models 4 and 5 are in agreement with Feldstein's statement that the economy was not in a recession in October or November 2007 (Feldstein, December 2007)

time series Y_t^* at time t as:

$$Y_t^* = \tilde{X}_t + \sum_{s=0}^{t-K-1} \theta^s \sigma(t-s) \varepsilon_{t-s}.$$

Thus, the latent time series conditional on the yield is multivariate normal with mean vector $[\tilde{X}_{t+1}, \tilde{X}_t, \tilde{X}_{t-1}]$ and variance matrix:

$$\begin{bmatrix} V(\tilde{Y}_{t+1}) & \theta V(\tilde{Y}_t) & \theta^2 V(\tilde{Y}_{t-1}) \\ \theta V(\tilde{Y}_t) & V(\tilde{Y}_t) & \theta V(\tilde{Y}_{t-1}) \\ \theta^2 V(\tilde{Y}_{t-1}) & \theta V(\tilde{Y}_{t-1}) & V(\tilde{Y}_{t-1}) \end{bmatrix}.$$

The results are then based on standard relationships between joint normals and conditional normals.

References

- [1] Burns, A. and W. Mitchell (1946), *Measuring Business Cycles*, New York: National Bureau of Economic Research.
- [2] Burns, A. (1960), "Progress Towards Economic Stability," *The American Economic Review*, Vol. 50, No. 1, 1-19.
- [3] Chauvet, M. (1998), "An Econometric Characterization of Business Cycle Dynamics with Factor Structure and Regime Switches," *International Economic Review*, Vol. 39, No 4, pp. 969-96.
- [4] Chauvet, M. and S. Potter (2001), "Recent Changes in the U.S. Business Cycle," *The Manchester School*, 69, 5, 481-508.
- [5] Chauvet, M. and S. Potter (2002). "Predicting Recessions: Evidence from the Yield Curve in the Presence of Structural Breaks," *Economics Letters*, Vol. 77, No. 2, 245-253.
- [6] Chauvet, M. and S. Potter (2005), "Forecasting Recessions Using the Yield Curve," *Journal of Forecasting*, 24, 2, 77-103.
- [7] Chauvet, M. and J.D. Hamilton (2006), "Dating Business Cycle Turning Points in Real Time," *Nonlinear Time Series Analysis of Business Cycles*, ed. Van Dijk, Milas, and Rothman, Elseviers Contributions to Economic Analysis series, 1-54.
- [8] Chauvet, M. and G. Popli (2008), "The Great Moderation Revisited: International Evidence," mimeo, University of California Riverside.
- [9] Chib, S., (2001), "Markov Chain Monte Carlo Methods: Computation and Inference" in J.J. Heckman and E. Leamer (eds.), *Handbook of Econometrics*, vol. 5, Amsterdam: North Holland, in press.
- [10] Dueker, M.J. (1999), "Conditional Heteroscedasticity in Qualitative Response Models of Time Series: A Gibbs Sampling Approach to the Bank Prime Rate." *Journal of Business and Economic Statistics*, 17, 466-72.
- [11] Dueker, M.J. (2005), "Dynamic Forecasts of Qualitative Variables: A Qual VAR Model of U.S. Recessions," *Journal of Business and Economic Statistics*, January, 23(1), 96-104.
- [12] Feldstein, M. (2007), "Wait Till Next Year," New York Times, December 16, 2007.
- [13] Geweke, J., (1999), "Using Simulation Methods for Bayesian Econometric Models: Inference, Development and Communication," *Econometric Reviews*, 18, 1-126.

- [14] Jeffrey, H. (1961), *Theory of Probability* 3rd ed., Oxford: Clarendon Press.
- [15] Koop, G. and S. Potter (1999), "Bayes Factors and Nonlinearity: Evidence from Economic Time series," *Journal of Econometrics*, 88, 251-281.
- [16] Koop, G. and S. Potter (2000), "Nonlinearity, Structural Breaks or Outliers in Economic time series?" in "*Nonlinearity and Econometrics* " edited by William Barnett et al., Cambridge University Press, May, 61-78.
- [17] Koop, G. and S. Potter (2007), "Estimating and Forecasting in Models with Multiple Breaks," *Review of Economic Studies*, July, 74, 763-789.
- [18] Potter, S. (2000), "A Nonlinear Model of the Business Cycle," in *Studies In Non-linear Dynamics and Econometrics*, July, 4, 2, 85-93.
- [19] McConnell, M. and G. Perez-Quiros (2000), "Output Fluctuations in the United States: What Has Changed Since the Early 1980s?" *American Economic Review*, 90, 5, 1464-1476.
- [20] Sensier, M. and D. van Dijk (2004) "Testing for Volatility Changes in US Macroeconomic Time Series, *Review of Economics and Statistics*, 86, 833-839.
- [21] Yates, J.F. (1982), "External Correspondence: Decompositions of the Mean Probability Score," *Organizational Behavior and Human Performance*, 30,132-156.

Table 1 - Posterior Mean Parameters across Models

Parameters	Model 1	Model 2		Model 3		Model 4	Model 5
		t<1984:01	t≥ 1984:01	τ < break	τ ≥ break		
β_0	-0.98	-0.56	-1.69	-0.63	-1.71	-0.67	-0.21
β_{IP}	-0.64	-0.53	-1.73	-0.45	-1.58	-0.52	-0.43
β_{Sale}	-0.17	-0.10	-0.32	-0.15	-0.23	-0.09	-0.37
β_{Income}	-1.18	-0.94	-2.16	-1.10	-2.21	-0.89	-1.05
$\beta_{Employment}$	-0.93	-0.33	-2.57	-0.71	-1.89	-0.89	-0.89
θ	-	-	-	-	-	-	0.70
Innovation Variance	1	1	1	1	1	-	-
1961:3-1970:11	-	-	-	-	-	1.09	1.84
1970:12-1975:3	-	-	-	-	-	0.54	0.63
1975:4-1980:7	-	-	-	-	-	3.84	2.54
1980:8-1982:11	-	-	-	-	-	0.23	0.54
1982:12-1991:3	-	-	-	-	-	0.21	0.23
1991:4-2001:12	-	-	-	-	-	0.15	0.19

Table 2 – Classification Table: Predictive Ability to Signal Recessions and Expansions

	Model 1			Model 2			Model 3			Model 4			Model 5		
	$Y_t=0$	$Y_t=1$	Total	$Y_t=0$	$Y_t=1$	Total	$Y_t=0$	$Y_t=1$	Total	$Y_t=0$	$Y_t=1$	Total	$Y_t=0$	$Y_t=1$	Total
$P(Y_t=1) \leq 0.5$	494	58	552	488	49	537	491	45	536	501	24	525	500	13	513
$P(Y_t=1) > 0.5$	9	24	33	15	33	48	12	37	49	2	58	60	3	69	72
Total	503	82	585	503	82	585	503	82	585	503	82	585	503	82	585
Correct	494	24	518	488	33	521	491	37	528	501	58	559	500	69	569
% Correct	98.2	29.3	88.5	97.0	40.2	89.1	97.6	45.1	90.3	99.6	70.7	95.6	99.4	84.1	97.3
% Incorrect	1.8	70.7	11.4	3.0	59.8	10.9	2.4	54.9	9.7	0.4	29.3	4.4	0.6	15.8	2.7
Total Gain	-1.8	29.3	2.6	-3.0	40.2	3.1	-2.4	45.1	4.3	-0.4	70.7	9.6	-0.6	84.1	11.3
% Gain	-	29.3	18.3	-	40.2	22.0	-	45.1	30.5	-	70.7	68.3	-	84.4	80.5

	Naïve Model		
	$Y_t=0$	$Y_t=1$	Total
$P(Y_t=1) \leq 0.5$	503	82	585
$P(Y_t=1) > 0.5$	0	0	0
Total	503	82	585
Correct	503	0	503
% Correct	100	0	86.0
% Incorrect	0	100	14.0

(*) The Total and % Gain is over the Naïve Model.

Recall that $Y_t=0$ for expansions and $Y_t=1$ for recessions. For example, model 1 correctly signals 24 out 82 recession observations for a probability cutoff of 50% $P(Y_t=1) > 0.5$.

Table 3 – Peak Signals of NBER Recessions

Peaks NBER	Model 1	Model 2	Model 3	Model 4	Model 5
1960:04	+1	+1	+1	+1	-3
1969:12	-2	-2	-2	-2	-2
1973:11	-3	+2	+2	-2	-3
1980:01	-3	-3	-3	-2	-4
1981:07	-5	-3	-3	-3	-3
1990:07	-3	-3	-3	-2	-4
2001:03	-2	-2	-2	-1	-2

(*) The criterion adopted to determine turning points is if the probability of recession is greater than 50%, $P(Y_t = 1) > 0.5$.

(**) Leads are represented by (+) and lags by (-). For example, model 5 indicates the beginning of the 2001 recession with a lag of two months.

Table 4 – Yates' Decomposition

	MSE	Var (x)	Δ Var (f)	Min Var (f)	$(\mu_f - \mu_x)^2$	2*cov (f, x)
Model 1	0.08410	0.12073	0.03398	0.00318	0.00001	0.07380
Model 2	0.08011	0.12073	0.04504	0.00660	0.00000	0.09226
Model 3	0.07820	0.12073	0.04380	0.00652	0.00013	0.09298
Model 4	0.05543	0.12073	0.05258	0.01643	0.00042	0.13474
Model 5	0.07075	0.12073	0.08329	0.04643	0.00027	0.17997

Yates' decomposition is: $MSE = Var(x) + \Delta Var(f) + Min Var(f) + (\mu_f - \mu_x)^2 - 2Cov(f, x)$, where x is the NBER dummy, f is the prediction from the model, var is the variance, μ is the mean, cov is the covariance, $Min Var(f) = (\mu_{f|x=1} - \mu_{f|x=0})^2 Var(x)$, and $\Delta Var(f) = Var(f) - Min Var(f)$.

Figure 1 – Posterior Distribution of Probability of a Breakpoint From Model 3




Figure 2 – Posterior Mean Probabilities of Recession for the Full Sample and NBER-Dating (Shaded Area)




Figure 3 – Posterior Mean Probabilities of Recession for the Last Two Decades and NBER Dated Recessions (Shaded Area)

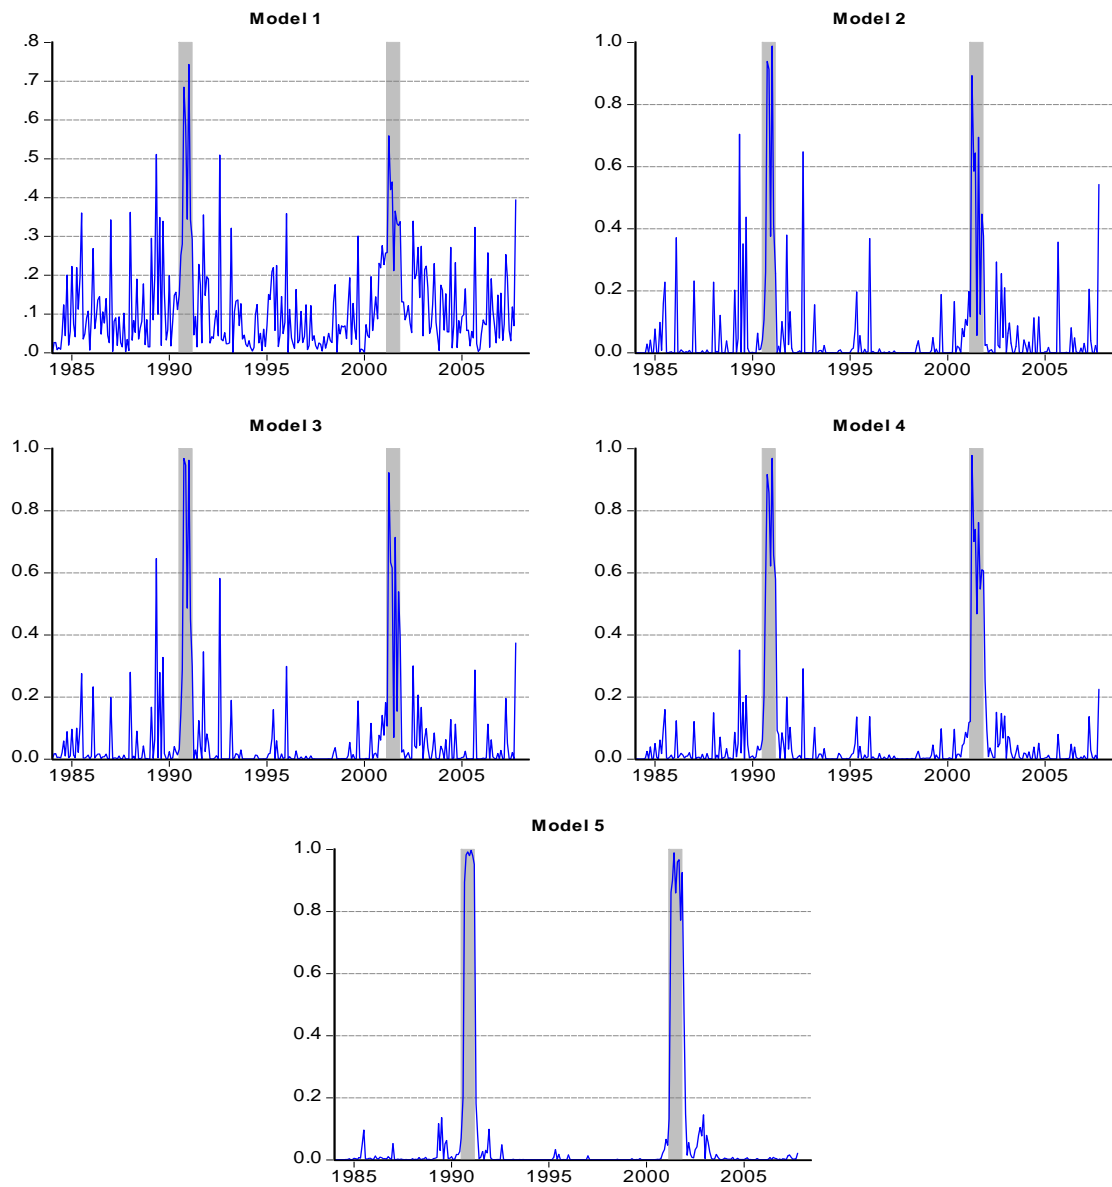


Figure 4 – Posterior Cumulative Distribution Function of the Probability of Recession in October 2007
Model 2 **Model 4**

