

Forecasting Recessions Using the Yield Curve*

Marcelle Chauvet and Simon Potter[†]

June 2001

Abstract

We compare forecasts of recessions using four different specifications of the probit model: a time invariant conditionally independent version; a business cycle specific conditionally independent model; a time invariant probit with autocorrelated errors; and a business cycle specific probit with autocorrelated errors.

The more sophisticated versions of the model take into account some of the potential underlying causes of the documented predictive instability of the yield curve. We find strong evidence in favor of the more sophisticated specification, which allows for multiple break-points across business cycles and autocorrelation. We also develop a new approach to the construction of real time forecasting of recession probabilities.

Keywords: Recession Forecast, Yield Curve, Structural Breaks, Bayesian, Classical Methods.

*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, CA, 92521. Tel: (909) 787-5037 x1587; email: chauvet@mail.ucr.edu. Address second author: Domestic Research, Federal Reserve Bank of New York, NY, 10045. Tel: (212) 720-6309; email:simon.potter@ny.frb.org.

1 Introduction

A large recent literature has shown that asset prices have significant predictive content for future economic activity.¹ In particular, recent empirical work has found evidence of systematic movements in the yield curve and future real output growth or recessions across a number of countries²

In general, the yield curve is upward sloped since long-term rates are higher than short-term ones. However, the slope of the curve tends to become flat or inverted before NBER-dated recessions. One of the possible reasons is that a tight monetary policy may precede a recession. Further, long-term rates reflect financial markets' expectation for future short-term rates. Hence, a flat or inverted curve might indicate that the market expects that future real interest rates should fall in face of a potential future recession or weak economic activity.

Although the yield curve is a statistically significant predictor of future activity, the predictive power of the spread is not stable over time. In particular, most models using the yield curve found it difficult to signal the 1990 recession in real time.³ One of the possible reasons for parameter instability of models using the yield curve is that its predictive power may depend on whether the economy is responding to real or monetary shocks (and implicitly on the monetary policy reaction function see Hamilton and Kim 2000 for a decomposition of the spread). In addition, there have been numerous changes in the market for U.S. Treasury debt over the last two decades. For example, recently the Treasury has been buying back debt leading to a reduction in the supply of long-term Treasury bonds.

Another potential reason is the recent changes in the volatility of the U.S. economy. In particular, McConnell and Perez (2000), Kim and Nelson (1999), and Chauvet and Potter (2000b) find evidence of a break towards more stability in the economy since 1984. Since these factors might affect the relationship between the yield curve and economic activity, models that do not take into account these evolving dynamics may lead to poor real time forecasts.

¹See Stock and Watson (2000) for a comprehensive literature review.

²The terms 'yield curve', 'term structure of interest rates', or simply 'term spread' refer to the difference in return between long term and short term government bonds.

³The exception is Laurent (1989). Harvey (1989) and Stock and Watson (1989) signaled the economic slowdown that started in 1989. See Stock and Watson's (2000) for a more detailed discussion.

In general, linear regression models that use output growth as the dependent variable indicate that the forecasting ability of the term spread has reduced since mid 1980s.⁴ However, models that focus instead on predicting a binary indicator of recession or expansion are more successful and stable over time than continuous ones.⁵ This evidence is corroborated in the recent work by Estrella, Rodrigues, and Schich (2000) (ERS hereafter), who examine the stability of the predictive model using classical tests for an unknown single breakpoint. They find evidence of structural break in the continuous model of U.S. industrial output in late 1983, but no evidence of instability over the full sample using binary models.⁶

In contrast with the results of ERS, Chauvet and Potter (2001a) find overwhelming evidence of structural instability when the binary probit model is estimated using the Gibbs sampler. The probability of recessions is substantially affected by consideration of a breakpoint as well as its location. Although the model used considered only a single break, there was evidence of the presence of multiple breakpoints. One possible resolution of these differing results is the presence of autocorrelated errors in the probit model. ERS adjust their test statistics for the presence of autocorrelation and heteroscedasticity under the null hypothesis, whereas Chauvet and Potter (2001a) use Bayesian tests that are not amenable to such non-parametric adjustment techniques.⁷ The potential effects on forecasting of not explicitly modeling correlation in the errors is large, and examining its importance is one of the main goals of this paper.

We extend the probit specification of Estrella and Mishkin (1998) (EM hereafter) to account for these two forms of potential misspecifications: time-varying parameters due to existence of multiple breakpoints, and the presence of autocorrelated errors. We examine the predictive content of the term

⁴See, for example, Haubrich and Dombrosky (1996), Dotsey (1998), Friedman and Kuttner's (1998) or Stock and Watson's (2000) survey.

⁵See Neftci (1996), Dueker (1997), Estrella and Mishkin (1998), and Estrella, Rodrigues, and Schich (2000).

⁶They use real industrial production and the spread between 10 year and 1 year interest rates from 1967:01 to 1998:12.

⁷For example, if we assume a break in the first month of 1984, the maximized likelihood of the model with a break is 1000 times that of a model without a break. In the classical test used by ERS this difference is not statistically significant after allowing for autocorrelated errors under the null hypothesis and the endogeneity of the breakpoint. In contrast, the Bayesian test compares the average (over the prior distribution) height of the likelihood between the two models and makes no correction for misspecification of either model.

structure in forecasting recessions under four different specifications: a time invariant conditionally independent version, as in EM, a time-varying conditionally independent version that takes into account multiple breakpoints across business cycles, a time invariant probit model with autocorrelated latent variable, and a time-varying probit model with autocorrelated latent variable.

In the standard probit model of EM conditional on the observed yield, the probabilities of the recession states are independent of each other (this follows directly from the assumption of independent errors). Figure 2a shows the implied probability of a recession state each month from 1955 through the end of 2000.⁸ Notice that until 1995-96 the probability of any particular month being a recession state was relatively low in the current expansion. Under the independence assumption, one can easily calculate the probability of not observing a recession over the 117 months from April 1991 to December 2000. For example, if the probability of a recession had remained fixed at the low value of 0.025, the probability of not observing a recession (i.e., continuing expansion for 117 months) would be $(0.975)^{117}$, which is approximately 5%. In other words, the probability of a recession occurring would be 95% (i.e., the probability that the first hitting time to a recession is 117 months or less). Alternatively, if we use the 60 estimated probabilities since the end of 1995 to 2000 from Figure 2a, we obtain a 12.5% probability of no recession through the period of 1996-1997, 5% through the period of 1996-1998, 0.5% through the period of 1996-1999, and an effectively 0 probability of no recession through the period of 1996-2000.⁹ These low probabilities of continuing expansion indicate that either the 1990s represent a period of extraordinary good luck or that the probit model suffers from some severe misspecification.

In this paper we estimate and compare probabilities of recession using alternative probit specifications. We use probabilities of continuing expansion and first recession time as described in the previous paragraph for producing forecasts. We believe that these probabilities provide more accurate real time

⁸The probabilities come from a standard probit model estimated by the Gibbs sampler (maximum likelihood techniques produce virtually identical results), using yield data from January 1954 to December 1999 and NBER business cycle dates from January 1955 to December 2000. Thus, we are assuming that the U.S economy was still expanding in December 2000.

⁹Note that the decrease in the probabilities of no recession around 1997 and 1999 is associated with the Asian and Russian currency crisis, respectively.

predictions of the likelihood of turning points.¹⁰ The models are estimated and analyzed using Bayesian methods. There are two motivations for the use of Bayesian techniques. The first is that maximum likelihood estimation for probit models with an autoregressive component is impractical for the size of dataset used.¹¹ In contrast, the Gibbs sampler can be used to simulate the latent variable of the probit model, which drastically reduces the computational complexity. Second, we are interested in presenting distributions over the forecast probabilities that contain information on both parameter uncertainty and uncertainty over the most recent value of the latent variable. These can be directly obtained from the output of the Gibbs sampler. In addition, it is computationally simple to calculate Bayes factors to compare the various models using the Savage-Dickey Density ratio within the Gibbs sampler.

The Bayes factor indicates overwhelming evidence in favor of the more complicated specifications. Plots of the in-sample predicted recession probabilities indicate that specifications with autocorrelated errors provide a much cleaner classification of the business cycle into expansion and recession periods. We also compare the performance of the alternative models in a real-time forecasting exercise using yield curve data from January 2000 to March 2001. While all the models considered indicate that the yield curve is currently signaling weak future economic activity in 2001-2002, we find that the strength of a recession signal differs substantially across specifications with weaker signals coming from the more complicated specifications.

The paper is organized as follows. The next section introduces the probit model and discusses our various extensions. The third section outlines the Gibbs sampler and the construction of the Bayes factors. In the fourth section the empirical results from the alternative specifications are discussed and some real time forecasting results presented. The fifth section concludes.

2 Extending the Probit Model

The probit model assumes an underlying latent variable Y_t^* for which there exists a dichotomous realization of an indicator Y_t denoting the occurrence or

¹⁰These probabilities consider uncertainty on both the parameters and the most recent values of the latent variable.

¹¹The difficulty comes from the fact that it is necessary to evaluate multiple integration over the unobserved lagged variable.

non-occurrence of an event. Here we assume that the variable Y_t^* represents the state of the economy and we use as its indicator the recession dating provided by the NBER. Y_t can take on two values, 0 if the observation is an expansion or 1 if it is a recession:

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

In addition, it is useful to have a notation for the specific dates of the business cycles. We assume that a business cycle starts the month after a NBER trough and continues up to the month of the next NBER trough. We denote the business cycle dates n by $t_{n-1} + 1, \dots, t_n$. Further, we indicate the dates of the expansion and recession phases of business cycle n by the sets E_n and R_n , respectively. Hence, the dates of business cycle expansions and recessions are given by $E = \cup E_n$ and $R = \cup R_n$, respectively.

The unobservable variable Y_t^* is related to the yield curve according to following regression:

$$Y_{t+K}^* = \beta_0 + \beta_1 S_t + \varepsilon_t, \quad (2)$$

where S_t is the spread between the 10-year and 3-month Treasury Bill rates, K is the forecast horizon of the latent variable in months, and β_i ($i = 0, 1$) are the regression coefficients. The error term ε_t is assumed to be independently distributed over time with a standard normal distribution. Notice that multiplying Y_{t+K}^* by any positive constant does not change the indicator variable Y_{t+K} . This implies that the coefficients β_i 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 in order to fix the scale of Y_{t+K}^* .

From equations (1) and (2) we obtain:

$$P(Y_{t+K}^* \geq 0 | S_t, \beta) = \Phi[\beta_0 + \beta_1 S_t], \quad (3)$$

where Φ is the cumulative distribution function of the standard normal distribution, and $P(Y_{t+K}^* \geq 0 | S_t, \beta)$ is the conditional probability of a recession at the forecast horizon K . Note that this is not necessarily the probability of most interest in forecasting.

As an alternative way of constructing probabilities considers the first hitting time to a recession given by:

$$H_R(t) = \{H : Y_{t+H}^* > 0, Y_{t+H-1}^* < 0, \dots, Y_{t+1}^* < 0\},$$

with the associated probability:

$$\pi_R(k, t) = P[H_R(t) = k] = P[Y_{t+k}^* > 0 | Y_{t+k-1}^* < 0, \dots, Y_{t+1}^* < 0](1 - \pi_R(k-1, t)).$$

Conditional on a sequence of values for $S_{t-K}^k = \{S_{t-K+1}, S_{t-K+2}, \dots, S_{t-K+k}\}$ we can evaluate these expressions by:

$$\pi_R(k, t) = \Phi[\beta_0 + \beta_1 S_{t-K+k}] \prod_{s=1}^{k-1} \{1 - \Phi[\beta_0 + \beta_1 S_{t-K+s}]\}. \quad (4)$$

Notice that this expression reflects a conditionally constant probability of recession, and it is directly related to the likelihood function of the observed data:

$$\ell(Y^T | S^{T-K}, \beta) = \prod_{t \in R} \Phi[\beta_0 + \beta_1 S_{t-K}] \prod_{t \in E} \{1 - \Phi[\beta_0 + \beta_1 S_{t-K}]\}, \quad (5)$$

where $\beta = [\beta_0, \beta_1]'$.

In addition, the ordering of $\prod_{s=1}^{k-1} \{1 - \Phi[\beta_0 + \beta_1 S_t]\}$ makes no difference to the hitting probability or to the value of the likelihood function. This is a direct consequence of the assumption of conditional independence and constant relationship between the yield spread and the probability of recessions. In particular, if we consider two expansions where the values of the term spread are given by permutations on the set

$$\{S_{t-K+1}, S_{t-K+2}, \dots, S_{t-K+k-1}\},$$

and fix the most recent value at S_{t-K+k} , they will have exactly the same probabilities of the expansion ending at time $T + k$.

We extend this framework in two main ways. First, we allow the variance of the innovation to change with the business cycle, which accounts for potential structural breaks in the relationship between the yield curve and the economy. Second, we add an autoregressive component to the model.

2.1 Business Cycle Specific Model

We consider a more general specification of the probit model in which the variance of the innovation may change across the N business cycles:

$$Y_t^* = \beta_0 + \beta_1 S_{t-K} + \sigma(t)\varepsilon_t, \quad (6)$$

where

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

and the initial business cycle is partially observed starting at $t = K + 1$ ($t_1 = K$). Hence, even in the case that the collection of yield spreads is the same, we can obtain different hitting probabilities across business cycle, since now we have:

$$\pi_R(t, k) = \Phi_n[\beta_0 + \beta_1 S_{t-K+k}] \prod_{s=1}^{k-1} \{1 - \Phi_n[\beta_0 + \beta_1 S_{t-K+s}]\}. \quad (7)$$

where

$$\Phi_n[\beta_0 + \beta_1 S_t] = \Phi[(\beta_0 + \beta_1 S_t)/\sigma_n].$$

There are two interpretations of this model. The first is the literal one that shocks to the latent variable may change over business cycles. For example, one could imagine that a long business cycle has a smaller innovation variance than a short one. The alternative interpretation arises from the fact that the scale of the innovation and the coefficient parameters β_i can not be separately identified. Thus, we can think of this as a time-varying parameter model, 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 $\beta_{n1} = \beta_1/\sigma_n$.

The likelihood function for the business-cycle specific variance is now:

$$\ell(Y^\top | S^{\top-K}, \beta, \{\sigma_n\}) = \prod_{n=1}^N \left(\frac{\prod_{t \in R_n} \Phi_n[\beta_0 + \beta_1 S_{t-K}]}{\prod_{t \in E_n} \{1 - \Phi_n[\beta_0 + \beta_1 S_{t-K}]\}} \right), \quad (8)$$

and now is grouped into particular expansions and contractions.

2.2 Dependence in the Latent Variable

In another generalization of the probit model we allow the latent variable Y_t^* to follow a first order autoregressive process. In this case we can no longer take permutations on the order of the yield spread over the previous $k - 1$ periods and obtain the same hitting probability. The latent variable model becomes:

$$Y_t^* = \beta_0 + \beta_1 S_t + \theta Y_{t-1}^* + \sigma(t)\varepsilon_t, n = 1, \dots, N. \quad (9)$$

where the autoregressive parameter $|\theta| < 1$. The hitting probabilities require one to integrate out over the unknown lagged value of the latent variable. Starting from an expansion state at time t in business cycle n we have:

$$\pi_R(t, 1) = \int_{y_t^* < 0} \Phi_n(\beta_0 + \beta_1 S_{t-k} + \theta y_t^*) f(y_t^* | S^t) dy_t^*,$$

where the complete sequence of yield spreads is part of the conditioning set, since it is now useful for inferring the unobserved value of Y_t^* . The k -period ahead hitting probability requires multiple integration over the values of $\{Y_{t+s}^* : s = 0, \dots, k - 1\}$:

$$\begin{aligned} \pi_R(t, k) = & \int_{y_t^* < 0} \cdots \int_{y_{t+k-1}^* < 0} \Phi_n[\beta_0 + \beta_1 S_t + \theta y_{t+k-1}^*] \\ & \prod_{s=0}^{k-1} \{1 - \Phi_n[\beta_0 + \beta_1 S_{t-k+s} + \theta y_{t+s}^*]\} \\ & f(y_{t+k-1}^*, \dots, y_t^* | S^t) dy_{t+k-1}^* \cdots dy_t^*. \end{aligned} \quad (10)$$

The likelihood function can then be written as the product of the recession hitting probabilities and the analogous expansion hitting probabilities, $\pi_E(t, k)$, using the appropriate business cycle dates and recession lengths r_n :

$$\begin{aligned} & \ell(Y^T | S^{T-k}, \beta, \{\sigma_n\}, \theta) \\ = & \pi_I(t_{N-1}, T - t_{N-1}) \prod_{n=2}^{N-1} \left\{ \frac{\pi_R(t_{n-1} + 1, t_n - t_{n-1} - r_n + 1)}{\pi_E(t_n - r_n, r_n)} \right\}, \end{aligned} \quad (11)$$

where $\pi_I(t_{N-1}, T - t_{N-1})$ is the probability of not observing a recession in the last $T - t_{N-1}$ periods of the sample, that is, the probability associated with the current continuing business cycle. This likelihood function now fully encodes the information on the lengths of particular expansions and contractions.

3 Estimation and Forecasting Techniques

The model with an autoregressive process for the latent variable is difficult to estimate using maximum likelihood, since multiple integration over the unobserved lagged variable is required. Thus, we use the Gibbs sampler to evaluate the likelihood function.¹² Further, since we are interested in out-of-sample estimates of the hitting probabilities of recessions, we focus on the properties of the joint posterior distribution of the model parameters and most recent value of the latent variable.

3.1 Obtaining Draws of the Latent Variable

The Gibbs sampler proceeds by generating draws of the latent Y_t^* conditional on $(\beta_0, \beta_1, \{\sigma_n\}, \theta)$ and the observed term spread. To simplify notation, let $X_t' \boldsymbol{\beta} = \beta_0 + \beta_1 S_{t-k}$. 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 presence of the lagged value of the latent variable in the conditional mean complicates the sampler. 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 $X_T' \boldsymbol{\beta} = \beta_0 + \beta_1 S_{T-k} + \theta Y_{T-1}^*$. This would generate a draw of the last period value of the latent variable.

Now 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_K^* . We simplify the analysis by assuming that $Y_K^* = \beta_0 + \beta_1 S_0 = 0$. Then, as shown in the appendix,

¹²See Geweke 1999 and Chib 2001 for excellent introductions to modern Bayesian computational techniques and references to earlier work on probit models that we implicitly draw on.

we have **a priori** the conditional normal distribution with mean:

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

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

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

and

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

Hence, we draw from the appropriate truncated normal as above to obtain a new draw of Y_{T-1}^* . This process continues until we arrive at the initial observation period. The value of Y_{K+1}^* is drawn in a similar manner to Y_T^* , by conditioning on the new draw of Y_{K+2}^* . However, this value has a different form, since its mean is given by:

$$X'_{K+1} \beta + \frac{\theta \sigma^2(1)}{\sigma^2(1) + \theta^2 \sigma^2(2)} \left(Y_{K+2}^* - \tilde{X}'_{K+2} \right),$$

and variance by:

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

3.2 Obtaining Draws of Probit Model Parameters

Given this sequence of draws for $\{Y_t^*\}$, we then need to obtain draws of the parameters $(\beta, \{\sigma_n\}, \theta)$. For simplicity, we use prior distributions from parametric families that generate simple conditional distributions for the posterior.¹³ We assume that the parameters of the conditional mean β are **a priori** bivariate normal with mean vector $\mu(\underline{\beta})$ and variance matrix $V(\underline{\beta})$. In addition, the autoregressive parameter θ is assumed to have an **a priori** truncated normal on $(-1, 1)$ with mean $\mu(\underline{\theta})$ and variance $V(\underline{\theta})$ independent 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 $Z_t = (Y_t^* - \theta Y_{t-1}^*)$. Then, conditional on $(\{Y_t^*\}, \{\sigma_n\}, \theta)$, the parameters β are obtained from a normal distribution with variance matrix:

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

and mean vector:

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

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

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

and mean:

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

¹³The sampler will require draws from normal and gamma distributions. Random draws from both types of distributions are available in widely used packages such as Gauss and Matlab.

If this draw satisfies the stationary condition, it is accepted. If not, it is rejected and new draws are made until one is obtained that satisfies the stationarity condition.

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 business cycle $n \geq 2$ by:

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

3.3 Bayes Factors

We are interested in the recession probabilities obtained from four different models:

1. Model 1: a probit specification with constant variance and serially uncorrelated latent variable;
2. Model 2: a probit specification with business cycle specific variance and serially uncorrelated latent variable;
3. Model 3: a probit specification with constant variance and autoregressive latent process;
4. Model 4: a probit specification with business cycle specific variance and autoregressive latent process.

We use Bayes factors to assess the value-added of the business cycle specific variances and autoregressive structure specifications. That is, we compare the marginal likelihoods of the various models, which correspond to the average height of the likelihood of the observed data (i.e., the NBER business cycle dates) with respect to the prior distribution. In some situations, the Bayes factor corresponds to the likelihood ratio statistic. For example, the likelihood ratio and Bayes factor are equal for the case of testing the probit model with autoregressive errors and time invariant variance, with $\theta = 0.8$, against the conditionally independent time invariant probit model with all other parameters known.

More generally, the Bayes factor differs from the likelihood ratio by the way it integrates out nuisance parameters and averages over possible parameter values for the alternative model. In order for these operations to make sense, it is necessary that the prior distributions are proper, that is, that they integrate to 1. For example, if we place an improper prior distribution on θ , its average height over the line is not well-defined. Further, even if we arbitrarily fixed a height for the prior distribution by construction, it would place zero weight on values of θ where the likelihood was higher than in the restricted model.

The major disadvantage of marginal likelihoods is their dependence on the prior distribution and the difficulty in their calculation. As explained above, diffuse prior distributions are not appropriate when calculating Bayes factors. However, we have prior distributions that are sufficiently uninformative so that the sample information will dominate. In our case, we have the general restriction that $\theta \in (-1, 1)$. In addition, for the business cycle specific variances model we choose a prior with a small number of degrees of freedom and with a mean equal to 1.

With respect to their calculation, the Bayes factor can be greatly simplified by using the Savage-Dickey Density ratio. If two models are nested, then there exists at least one point in the parameter space of the unrestricted model where its likelihood is equivalent to the restricted model. For example, if we evaluate our most general model at $\theta = 0$, then its likelihood is equivalent to the business cycle specific probit. In the nested case, the Bayes factor can be found by the ratio of posterior density at $\theta = 0$ to the prior density at $\theta = 0$, under some mild regularity conditions. Continuing with the example the Bayes factor would be:

$$\text{BF}_{2 \text{ vs } 4} = \frac{p(\theta = 0|\text{Data})}{p(\theta = 0)}.$$

Notice that with Gibbs sampling we do not have direct access to the unconditional posterior densities. Instead, we can calculate them at each iteration of the sampler as:

$$p(\theta = 0|\text{Data}, \mathbf{Y}^*, \beta, \{\sigma_n\}),$$

where $\mathbf{Y}^* = [Y_{K+1}^*, \dots, Y_T^*]'$. However, since

$$\begin{aligned} p(\theta = 0|\text{Data}) &= \int p(\theta = 0|\text{Data}, \mathbf{Y}^*, \beta, \{\sigma_n\}) p(\mathbf{Y}^*, \beta, \{\sigma_n\}|\text{Data}) d\mathbf{Y}^* d\beta d\sigma_2 \cdots d\sigma_n, \end{aligned}$$

we can average these values across draws of the sampler to obtain an estimate.

The Bayes factor (BF) is assessed using the recommendations of Jeffrey's (1961, Appendix B), which states the weight of evidence against the null given by the following range of values:

- $\ln BF > 0$ 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.

3.4 Real Time Forecasts

The main objective of the analysis of the relationship between the yield curve and business cycle turning points is to provide better real time predictions of turning points. In the original probit model of EM, forecasts were constructed by using the maximum likelihood estimates and recent observations on the yield curve. For example:

$$P(Y_{T+K}^* \geq 0 | S_T, \hat{\beta}) = \Phi[\hat{\beta}_0 + \hat{\beta}_1 S_T].$$

Note that since it is unlikely that the predicted probability is exactly equal to 0 or 1 some judgement is required on how to interpret the predictions. We have emphasized a different type of prediction in Section 2, namely the probability of an expansion continuing for k more months. For the original probit specification of EM, this forecast is obtained from:

$$\prod_{s=1}^k \left\{ 1 - \Phi[\hat{\beta}_0 + \hat{\beta}_1 S_{T+k-s}] \right\},$$

where this forecast converges to 0 as k increases by construction.

The first difference in prediction using the Gibbs sampler approach is that instead of evaluating the cumulative distribution function using the maximum likelihood estimator, one evaluates it at each draw of the Gibbs

sampler. The collection of forecasts can then be averaged to produce an estimated posterior mean probability of recession:

$$\hat{P}(Y_{T+K}^* \geq 0 | S_T, S^{T-K}, \psi, Y_T) = \frac{1}{I} \sum_{i=1}^I \Phi[\beta_0^{\{i\}} + \beta_1^{\{i\}} S_T],$$

where ψ signifies the hyperparameters of the prior distributions and the assumption on the initial conditions of the latent variable and the yield curve. Alternatively, the posterior predictive distribution can be analyzed directly and used to present probability intervals on the probability of recession. Analogously, one can form the cumulative product of the individual probabilities at each iteration of the Gibbs sampler and examine its posterior properties.

The second difference arises from the more complicated model with an autoregressive component. In this case, as noted above, one needs to integrate out over the unknown lagged value of the latent variable to form predictions. This requires a numerical integration step at each iteration of the Gibbs sampler. The Gibbs sampler produces a draw of Y_T^* , θ , and σ_N . These values can be used to simulate J time series realizations of $\{Y_{T+s}^*; s = 1, \dots, k\}$ using the observed yield values from $T - K + 1$ to $T - 1$. The average over these J draws is formed as:

$$\begin{aligned} \hat{P}(Y_{T+K}^* \geq 0 | S_T, S^{T-K}, \psi, Y_T^{*\{i\}}, \xi^{\{i\}}) = \\ \frac{1}{J} \sum_{j=1}^J \Phi \left[(\beta_0^{\{i\}} + \beta_1^{\{i\}} S_T + \theta^{\{i\}} Y_{T+K-1}^{*\{j\}}) / \sigma_N^{\{i\}} \right], \end{aligned}$$

where $\xi^{\{i\}}$ denotes the set of parameter draws at the i th iteration of the Gibbs sampler. Averaging this estimated probability over the I draws from the Gibbs sampler produces an estimate of the posterior mean of the probability of recession and the collection of draws provides an estimate of the posterior distribution.

4 Results

4.1 Priors

We assume that the prior variance of β is the identity matrix. For the prior mean we use the maximum likelihood estimate from model 1, which assumes

unit variance and no autoregressive component, as in EM. Given the size of the sample there is little influence from the prior. However, we decided to give to EM's standard probit model the advantage of a prior centered at its MLE. For the autoregressive parameter θ the prior is taken to be a truncated standard normal on $(-1, 1)$. With respect to the prior for the business cycle specific variances, we center it at 1, with diffuse degrees of freedom equal to 3.

4.2 Empirical Results

The data on the yield curve cover the period from January 1954 to March 2001.¹⁴ We set the forecast horizon K to 12 and assume that no recession will be called in the last three months of 2000.¹⁵ Thus, our business cycle data run from January 1955 to December 2000. The last 15 observations of the yield spread are used to form real time forecasts of the probability of recession state for each month from January 2001 to March 2002, and the probability of an expansion continuing through March 2002.

We use 40,000 iterations to estimate the probit models using the Gibbs sampler. We start the sampler from the maximum likelihood estimate for the whole sample from Model 1, but calculate the posterior properties only after 10,000 draws (thus, 50,000 draws in total).¹⁶

Table 1 summarizes the posterior means of the parameters for the four alternative specifications and also gives some forecasts of the probability of recession in 2001 – 2002. The results from all models corroborate previous findings, indicating a significant relationship between inversions of the term structure and the probability of a recession 12 months ahead.

The posterior means of the parameters for the different models can not be directly compared due to the presence of the autoregressive component and/or business cycle specific variance. However, we can observe that in

¹⁴The yield spread is defined as the difference between the 10-year constant maturity Treasury Bond and the 3-month constant maturity Treasury Bill. We take the average of the daily rates each month.

¹⁵Hall (2001) suggests that there the NBER is unlikely to call a recession in the last quarter of 2000. Robert Hall is one the members of the NBER Business Cycle Dating Committee.

¹⁶The computation time for the most complicated model is around 4 hours, including the construction of the out-of-sample forecasts with the use of 1000 simulated draws for each Gibbs draw.

all cases the posterior mean of the intercept and slope coefficients are negative as expected. Further, the autoregressive parameters in Models 3 and 4 are relatively large and positive, implying highly persistent movements in the underlying latent variable. These results affect the recession forecasts substantially, as discussed below.

Table 1 shows the different values obtained for the variance across the seven business cycles in the sample for Models 2 and 4, which allow the variance to change over the business cycle. Except for a scale factor, the estimated variances from these models capture the same features across different business cycles in the sample. In particular, the highest estimated variance occurs during the short 1980-81 business cycle. This corresponds to the period in which the Federal Reserve changed its operating procedures. The level and volatility of interest rates and consequently, of the yield curve, increased substantially during the 1979-82 period. On the other hand, the variances with lowest values are associated with the long expansions of the 1960s, 1980s and 1990s.

The negative of the posterior means for the latent variable (Y_t^*) are plotted in Figure 1 together with the yield spread for all four models considered. We can observe that the posterior mean of the latent variable from the simplest version of the probit specification (Model 1) does not vary much as compared to the other models, and essentially tracks the yield curve.

This can also be observed in Figure 2, which plots the posterior mean of the probabilities of recession from January 1955 to March 2002 for all four specifications. The probabilities consistently rise before each of the seven full recessions in the sample dated by the NBER. However, the probabilities of a recession for Model 1 are only above 50% for the 1974-75, the 1980, and the 1981-1982 recessions. In addition, Model 1 gives more false signals of recessions than the other specifications. Notice that the probabilities of recessions from this model rise as much for recessions as for the low-growth phases in the U.S. economy in 1966-67, in 1995, and in 1997, which are associated, respectively, with slowdowns in Europe, with the Mexican crisis, and with the East Asian crisis. As a result, interpretation of increases in probabilities as an indication of future recessions is not unambiguous. This is also the case for the current slowdown in 2000-2001, as discussed in the next section.

Model 2 allows for the possibility of time-varying parameters or, more specifically, for variances that may change for each of the recessions in the sample. There are two main differences between the posterior probabilities

of recession from Models 1 and 2. First, the posterior probabilities rise substantially more before each of the NBER-dated recessions in Model 2 than in Model 1. Second, Model 2 gives only one false signal of recession – the probabilities forecast the 1966-67 low-growth phase as a recession. Thus, probabilities of recession from Model 2 are easier to interpret as signalling turning points. An interesting feature of this model is that the posterior latent variable increases substantially after the 1990-91 recession. Perhaps the magnitude of the increase in the latent variable is associated with the long extension of the 1990s expansion.

When allowing for an autoregressive process in the probit model, the estimated parameters and posterior probabilities of recessions are quite different from the other specifications. As seen in Figure 2, the posterior probabilities from Model 3 and 4 are very similar to each other and very different from Models 1 and 2. In particular, the probabilities of recession from Models 3 and 4 consistently increase above 80% before each of the recessions in the sample. Thus, there is much less uncertainty regarding interpretation of these probabilities. In addition, the probabilities increase only before recessions, not before slowdowns. That is, the probabilities from these models do not give any false signals of recession.

The Bayes factor allows us to evaluate the sample evidence in favor of the alternative models. The Bayes factor indicates overwhelming evidence in favor of the specification that includes both business cycle specific variance and an autoregressive process for the latent variable. This could be expected given the significant improvement in fit produced by the extensions to the basic probit model, as exhibited in Figure 2. Further, the Bayes factor always favors the more complicated model in each possible pairwise comparison. First, the natural logarithm of the Bayes factor for constant variance versus business cycle specific variance is -55 (for Model 1 versus 2) and -27 (for Model 3 versus 4) showing decisive evidence for the time varying variance. Second, the Bayes factor for no autoregressive term versus an autoregressive term is impossible to distinguish from computer zero. Thus, the sample evidence decisively supports Model 4.

4.3 Recession Forecasts for 2001-2002

In this section, we illustrate the differences between the various models and also between the simple recession probabilities and hitting time probabilities in a real time forecasting exercise. The exercise uses the 15 months of yield

spread data after the end of the estimation sample in December 2000

Figure 3 plots the posterior mean of probabilities of recessions for 2001-2002 for all models. As it can be observed, the probability of recessions for this period differ considerably across the alternative specifications. In particular, the simplest versions of the probit model that do not correct for serially correlated errors have the highest probabilities of recession states in 2001-2002 (Models 1 and 2). Model 2 indicates that the U.S. economy may experience a short recession during the second semester of 2001, while Model 1 gives more moderate signals of a weak economy. On the other hand, the more sophisticated specifications that allow the latent variable to be serially correlated yield substantially smaller probabilities of a recession in the near future (Models 3 and 4).

More specifically, the probabilities of a recession state for particular months in the second semester of 2001 for Model 1 are around 40%. They reach a peak of 46% in December 2001 and decrease steadily to 27% in March 2002, the last month in the forecast exercise. However, one of the problems with Model 1, in addition to be misspecified, is that it equally signals both recessions and slowdowns. Thus, the probabilities may be indicating either a potential recession or simply the continuation of the low-growth phase in the U.S. economy that started in 2000. For Model 2, the probabilities of a recession in the second semester of 2001 are substantially higher (around 59%) reaching a peak in December 2001 of 70%. Given the recession probability history from this model, as observed in Figure 1b, these high values indicate that the economy may experience a short recession in the second semester of 2001. However, the probabilities of recession decline quickly to only 11% in March 2002.

A more accurate reading on the likelihood of recessions can be obtained from the hitting probabilities. Some information on these is shown in the last two rows of Table 1 and in Figures 4. The results are similar for Models 1 and 2, as the probability of continued expansion in March 2002 is less than 1% for both models. Figures 4a and 4b plot the posterior cumulative distribution of the hitting probabilities of no recession before March 2002 for these models. As seen in the figures, under the assumption of no autoregressive component in the latent variable 95% of the posterior on the probability of continued expansion is between 0.1% and 0.4% for Model 1, and between 0% and 0.9% for Model 2. Finally, Figures 5a and 5b show the cumulative distribution function of the probability of a recession state in March 2002. As observed, 95% of the posterior on the probability of a recession is between 0 and 28%

indicating that the positive steepening of the yield curve in early 2001 signals stronger economic activity in 2002.

On the other hand, when we allow the latent variable to follow an autoregressive process, the resulting recession probability forecasts are substantially smaller. In particular, the probability of a recession in the second semester of 2001 is around 18% for both models 3 and 4 (Figures 3c and 3d). The probabilities reach a maximum in January 2002 of only 25% for Model 3 and 31% for Model 4, and slowly decrease until March 2002. As observed in Figures 2c and 2d, the probabilities from these models increase above 80% before each of the recessions in the sample.

However, the evidence from the hitting probabilities show higher uncertainty with respect to a future recession. In particular, the probabilities that an expansion will continue in March 2002 are around 50% for Models 3 and 4 (Table 1). As illustrated in Figures 4c, 4d, 5c and 5d, if the model explicitly addresses the error misspecification by introducing an autoregressive process, the inferences of recession probabilities as well the forecasts of recession in March 2002 change considerably. Figures 5c and 5d plot the posterior distribution of the probability of a recession state in March 2002 for these models. While the estimated mean recession probability for Model 3 is 24%, a symmetric probability interval of 95% around this mean is around 2% and 51%. For Model 4, this symmetric probability interval around the mean recession probability of 31% in March 2002 is between 0% to 64%. That is, the uncertainty with respect to a recession probability increases somewhat when taking into account multiple breakpoints and serial correlation in the conditional mean of the term structure. Further, Figures 4c and 4d show the cumulative distribution function of the probability of no recession before March 2002 for Models 3 and 4. We observe that 95% of the posterior on the probability of no recession before March 2002 is between 12% and 89% for Model 3, and between 0.6% and 96% for Model 4.

In summary, the posterior recession probabilities for all four models peak around either December 2001 or January 2002. Further, all models indicate a subsequent smaller probability of recession state in 2002 with the reduction in risk lower for the models with autocorrelation. However, the magnitude of the probabilities differ substantially across model specifications.

4.4 Previous Forecast Performance

We consider out-of-sample predictions for four periods, which were estimated using data up to the month preceding the starting date of the forecast intervals below:¹⁷

1. December 1988 to February 1990 using yield data up to February 1989,
2. December 1989 to February 1991 using yield data up to February 1990,
3. January 1999 to March 2000 using yield data up to March 1999,
4. January 2000 to March 2001 using yield data up to March 2000.

Notice that the first and third periods do not contain a recession, while the second period includes the 1990-91 recession. It is possible that the fourth period might contain a recession in its last two months.

Table 2 contains some posterior features of this modified out-of-sample forecasting exercise implemented for Model 4. The quantities in parentheses represent the analogous results obtained using the original probit model of EM (Model 1). The first thing to note is that the posterior means of the parameters are relatively stable across the estimation periods. The only minor exception is when the data from January 1999 to December 2000 is incorporated (i.e., moving from period 3 to period 4).

With respect to the hitting probabilities, the probability of continuing expansion from January 1999 to March 2000 is relatively low at 0.66. Although it is providing a stronger recession signal than period 2, it is not as strong as the current prediction of 0.49 (Table 1). From Figure 1 we can observe that the term spread was relatively low at the end of 1998, but the lack of a subsequent recession in 1999 leads to an attenuation of the effect of the spread compared to the constant term. On the other hand, the importance of the spread increases with the introduction of the 1990s into the sample, indicating that the large positive spreads for most of the period were consistent with the long expansion observed¹⁸.

¹⁷EM find that the in-sample results for the standard probit specification did not necessarily match with the out-of-sample performance. For the more complicated versions considered here, the computational burden is too high to repeat the the comprehensive recursive forecasting exercise conducted by these authors.

¹⁸That is, the slope coefficient divided by the innovation standard deviation is substantially larger in absolute value.

The model is relatively successful in forecasting the 1990-91 recession in the sense that it did not produce a false signal in 1989. However, only a weak recession signal is produced before the actual recession, as found by previous models. In particular, the probability of continuing expansion for period 2 is 0.70. Notice that if we return to the recession forecast for the 1999 period, there was considerable speculation regarding the end of the expansion in the Fall of 1998, given the financial turmoil experienced by Russia and fears of financial contagion. These discussions were cut short with the strong growth of the US economy in the first quarter of 1999. Similar speculations were taking place before the 1990 recession, but the strong economic growth in the first quarter of 1990 also faded away discussions of a recession. However, in 1990 a recession hit the economy in the third quarter of the year.

The performance of the EM model is similar to the more complicated model in a qualitative sense. However, the standard probit model of EM gives hitting probabilities that are poorly calibrated. In particular, the posterior mean of the probability of continuing expansion from this model tend to be very low. Even for the prediction in period 1, which does not include a recession (but does include a slowdown), the upper 97.5% percentile of the probability of a continuing expansion is only 0.5. Thus, the model is basically signalling the 1989 slowdown. Given that generally slowdowns precede recessions, but not all slowdowns turn into recessions, this suggests that the more complicated probit model gives a more accurate quantitative evidence on the likelihood of recession.

5 Conclusions

This paper extends the probit specification of EM to account for the possibility of multiple breakpoints and serially correlated errors. We use the alternative specifications to construct hitting probabilities to the next recession. We find that a probit model with business cycle specific innovation variance and an autoregressive component has a much better in sample fit than the original probit model of EM. In particular, the recession forecasts from this more complicated model are very different from the ones obtained from the standard probit specification. In addition, the hitting probabilities suggest substantial misspecification in the standard probit model of EM.

All specifications considered indicate that the yield curve is signaling weak future economic activity in 2000-2001. However, the strength of the

recession signals differs substantially across the specifications. The versions of the probit model that do not correct for serially correlated errors display the highest posterior probability of a recession for 2001-2002. On the other hand, the more sophisticated specifications that allow the latent variable to be serially correlated lead to smaller probabilities of a recession in the near future.

Appendix

In this appendix we derive the formula required for the Gibbs sampler draws in the case of Model 4. First note that 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 $Y_{t+1}^*, Y_t^*, Y_{t-1}^*$ has a joint normal distribution, the conditional distribution is normal. Under the assumption that all initial values are zero, we can write the latent 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.

Table 1: Posterior Means Across Models

Estimated Parameters and Probabilities	Model 1	Model 2	Model 3	Model 4
β_0	-0.627	-0.342	-0.165	-0.099
β_1	-0.718	-1.041	-0.209	-0.228
θ	-	-	0.875	0.848
Innovation Variance				
1955-1957	1	1	1	1
1957-1961	-	4.086	-	1.899
1961-1970	-	0.191	-	0.266
1970-1975	-	1.499	-	1.663
1975-1980	-	0.793	-	0.716
1980-1982	-	45.542	-	3.315
1982-1991	-	0.780	-	0.476
1991-present	-	0.182	-	0.295
Probability of Continued Expansion through March 2002	0.009	0.003	0.517	0.493
Probability of Continued Expansion				
Lower 2.5th percentile	0.001	0.000	0.205	0.135
Upper 97.5th percentile	0.025	0.022	0.854	0.923

Table 2: Simulated Out of Sample Properties For Model 4

Estimated Parameters and Probabilities	Period 1	Period 2	Period 3	Period 4
β_0	-0.11 (-0.59)	-0.11 (-0.60)	-0.08 (-0.59)	-0.11 (-0.62)
β_1	-0.18 (-0.66)	-0.18 (-0.66)	-0.25 (-0.71)	-0.18 (-0.71)
θ	0.86	0.86	0.84	0.86
Innovation Variance				
1982-1991	0.44	0.44	0.44	0.51
1991-present	-	-	0.37	0.32
Probability of Continued Expansion Next 15 months	0.95 (0.38)	0.70 (0.02)	0.66 (0.04)	0.82 (0.17)

The numbers between parentheses correspond to the results from Model 1, for comparison.

References


- [1] Chauvet, M. and S. Potter (2001a), “Predicting a Recession: Evidence from the Yield Curve in the Presence of Structural Breaks,” **Working Paper**, University of California, Riverside.
- [2] Chauvet, M. and S. Potter (2001b), “Recent Changes in the U.S. Business Cycle,” forthcoming, **The Manchester School**.
- [3] Chib, S., 2001b, “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.
- [4] Dotsey, M. (1998), “The Predictive Content of the Interest Rate Term Spread for Future Economic Growth,” **Federal Reserve Bank of Richmond Economic Quarterly**, 84, 3, 31-51.
- [5] Dueker, M.J. (1997), “Strengthening the Case for the Yield Curve as a Predictor of U.S. Recessions,” **Federal Reserve Bank of St. Louis Review**, 79 (2), 41-50.

- [6] Estrella, A. and G. Hardouvelis (1991), "The Term Structure as a Predictor of Real Economic Activity," **Journal of Finance**, 46 (2), 555-576.
- [7] Estrella, A. and F. S. Mishkin (1998), "Predicting U.S. Recessions: Financial Variables as Leading Indicators," **The Review of Economics and Statistics**, 80, 45-61.
- [8] Estrella, A., A.P. Rodrigues, and S. Schich (2000), "How Stable Is the Predictive Power of the Yield Curve? Evidence from Germany and the United States," mimeo, Federal Reserve Bank of New York.
- [9] Friedman, B.M. and K.N. Kuttner (1998), "Indicator Properties of the Paper-Bill Spread: Lessons from Recent Experience," **The Review of Economics and Statistics**, 80, 34-44.
- [10] Geweke, J., (1999), "Using Simulation Methods for Bayesian Econometric Models: Inference, Development and Communication," **Econometric Reviews**, 18, 1-126.
- [11] Hamilton, J.D. and D.H. Kim, (2000), "A Re-Examination of the Predictability of Economic Activity Using the Yield Spread," mimeo, UCSD.
- [12] Harvey, C.R. (1989), "Forecasts of Economic Growth from the Bond and Stock Markets," **Financial Analysts Journal**, 45 (5), 38-45.
- [13] Haubrich, J.G. and A.M. Dombrosky (1996), "Predicting Real Growth Using the Yield Curve," **Federal Reserve Bank of Cleveland Economic Review**, 32 (1), 26-34.
- [14] Jeffrey, H. (1961), **Theory of Probability** 3rd ed., Oxford: Clarendon Press.
- [15] Kim, C-J. and C. Nelson (1999), "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov Switching Model of the Business Cycle," **Review of Economics and Statistics**, 81(4) pp 1-10.
- [16] Koop, G. and S. M. Potter (1999), "Bayes Factors and Nonlinearity: Evidence from Economic Time series," **Journal of Econometrics**, 88, 251-281.


- [17] Laurent, R. (1989), “Testing the Spread,” **Federal Reserve Bank of Chicago Economic Perspectives**, 13, 22-34.
- [18] 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.
- [19] Neftci, S. (1996): *An Introduction to the Mathematics of Financial Derivatives*, New York: Academic Press.
- [20] Stock, J. and M. Watson (1989), “New Indices of Coincident and Leading Indicators,” In O. Blanchard and S. Fischer edited **NBER Macroeconomic Annual** Cambridge, MIT Press.
- [21] Stock, J.H. and M.W. Watson (2000), “Forecasting Output and Inflation: The Role of Asset Prices,” *mimeo*, Kennedy School of Government, Harvard University.

Figure 1 - Negative Posterior Mean Latent Variable Under Different Assumptions (—), Actual Yield Curve (St-12) (---), and NBER Recessions (Shaded Area)


a) Model 1 - Constant Variance and i.i.d. Errors



b) Model 2 - Changing Variance and i.i.d. Errors



c) Model 3 - Constant Variance and AR(1) Process



d) Model 4 - Changing Variance and AR(1) Process






Figure 2 - Posterior Mean Probabilities of Recession Under Different Assumptions, and NBER Recessions (Shaded Area)


a) Model 1 - Constant Variance and i.i.d. Errors



b) Model 2 - Changing Variance and i.i.d. Errors



c) Model 3 - Constant Variance and AR(1) Process



d) Model 4 - Changing Variance and AR(1) Process






Figure 3 - Posterior Mean Forecasts of Recession Under Different Assumptions for 2001/2002


a) Model 1 - Constant Variance and i.i.d. Errors



b) Model 2 - Changing Variance and i.i.d. Errors



c) Model 3 - Constant Variance and AR(1) Process



d) Model 4 - Changing Variance and AR(1) Process




Figure 4 – Posterior Cumulative Distribution Function of the Probability of No Recession Before March 2002 Under Different Assumptions




Figure 5 – Posterior Cumulative Distribution Function of the Probability of Recession in March 2002 Under Different Assumptions

