

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

**Incorporating Vintage Differences and Forecasts into Markov
Switching Models**

Jeremy J. Nalewaik

2007-23

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Incorporating Vintage Differences and Forecasts into Markov Switching Models

Jeremy J. Nalewaik*

March 12, 2007

Abstract

This paper discusses extensions of standard Markov switching models that allow estimated probabilities to reflect parameter breaks at or close to the end of the sample, too close for standard maximum likelihood techniques to produce precise parameter estimates. The basic technique is a supplementary estimation procedure, bringing additional information to bear to estimate the statistical properties of the end-of-sample observations that behave differently from the rest. Empirical results using real-time data show that these techniques improve the ability of a Markov switching model based on GDP and GDI to recognize the start of the 2001 recession.

*Board of Governors of the Federal Reserve System (e-mail: jeremy.j.nalewaik@frb.gov). The views expressed in this paper are solely those of the author and not those of the Board of Governors of the Federal Reserve System. I thank seminar participants at the Board of Governors for comments.

1 Introduction

This paper discusses extensions of standard Markov switching models that allow estimated probabilities to reflect parameter breaks at or close to the end of the sample, too close for standard maximum likelihood techniques to produce precise parameter estimates. The basic technique is a supplementary estimation procedure, bringing additional information to bear to estimate the statistical properties of the end-of-sample observations that behave differently from the rest; the additional information required is the historical values of the end-of-sample observations over the entire length of the time series employed. For example, assume that the first official estimate of GDP growth for a given quarter behaves differently than later-vintage estimates that have been revised a number of times, so a differential treatment of these vintages is appropriate. The techniques discussed here require the first release of GDP for every period in the sample under consideration. While this may sound like an onerous requirement, databases such as the Federal Reserve Bank of Philadelphia's Real-Time Dataset make it increasingly simple to meet.

The supplementary estimation techniques proposed here are computationally simple and highly intuitive. Many economists have computed the means and variances of variables of interest in recessions and expansions, defined by the NBER's start and end dates. One can interpret such a mean in recession periods as a weighted mean, with a weight of one assigned to periods designated as in recession, and a weight of zero assigned to other periods. The supplementary estimation procedure does essentially the same thing, taking weighted means and variances of the variables of interest - the time series of end-of-sample observations - with the weights being probabilities of recession determined by first-stage estimation of a Markov switching model, so the weights are not constrained to be either zero and one as implied by the NBER dating. Hamilton

(1990) showed that such weighted averages are typically the maximum likelihood estimates of Markov switching model parameters; those results motivate the supplementary estimation procedure proposed here.

Section 2 of the paper discusses the crux of the basic problem in a general context, while section 3 outlines the solution the paper proposes. Section 4 works through three empirical applications in the context of estimating probabilities of recession in real time for the U.S. economy. The first application studies vintage differences in GDP growth, comparing the properties of the early-vintage estimates versus later-vintage estimates that have passed through annual and benchmark revisions. The supplementary estimation procedure employed has interesting implications, whether or not it is employed to modify standard Markov switching models: it serves as a useful diagnostic check on the early-vintage GDP estimates, asking whether they give useful information about the state of world.

The second application considers appending the officially-published GDP growth time series with forecasts of GDP growth. The forecasts used in this application are Greenbook forecasts produced by the staff of the Board of Governors of the Federal Reserve. These forecasts are released to the public with a lag of roughly five years, so this application is not meant to be of real-time use to anyone outside the Fed; it is merely meant to illustrate how a long history of forecasts may be used in conjunction with the paper's estimation techniques. The Board staff closely tracks the source data used to compute GDP growth, producing near-term GDP growth forecasts that are quite accurate. Updating probabilities of recession with these forecasts for the current quarter, instead of waiting to update until the official GDP data is released, has the potential to substantially increase the timeliness of the estimated probabilities.

The third application studies a bivariate Markov switching model estimated using

GDP and gross domestic income (GDI), following Nalewaik (2007). The application incorporates vintage differences and appends each quarterly time series with forecasts based on higher-frequency indicators such as industrial production or unemployment insurance claims. This application uses real time data to track the performance of the model around the start of the 2001 recession, measuring the incremental impact of including the forecasts in the model. The results confirm that the forecasts do improve the timeliness of the model in recognizing the start of that recession.

Section 5 discusses alternative procedures for incorporating vintage differences and forecasts into Markov switching models, and points out reasons why the approach described here may be preferable. Section 6 concludes.

2 The Basic Problem

The standard Markov switching model outlined in Hamilton (1994) describes the vector of n variables of interest \mathbf{Y}_t for periods $t = 1, \dots, T$. Reviewing the standard model briefly, the joint density of \mathbf{Y}_t conditional on the state of the world in the current period, S_t , is $f(\mathbf{Y}_t|S_t)$. Past states of the world, S_{t-k} , have no effect on this conditional density. There exist J states of the world in the model, and let the estimated probability of state j in period $t - 1$, conditional on history through period $t - 1$, be:

$$\text{prob}(S_{t-1} = j \mid \mathcal{H}_{t-1}) = \pi_{t-1=j|t-1},$$

where $\mathcal{H}_{t-1} = \{1, \mathbf{Y}_{t-1}, \mathbf{Y}_{t-2}, \dots, \mathbf{Y}_2, \mathbf{Y}_1\}$. Stack the J probabilities $\pi_{t-1=j|t-1}$ into a column vector $\xi_{t-1|t-1}$. The set of probabilities of state j occurring next period, conditional on state i today, can be gathered into a $J \times J$ Markov transition matrix P , with

typical element $p_{ij} = \text{prob}(S_t = j | S_{t-1} = i)$. Application of the transition matrix yields:

$$(1) \quad \xi_{\mathbf{t}|\mathbf{t}-1} = P\xi_{\mathbf{t}-1|\mathbf{t}-1}.$$

Given the J elements of $\xi_{\mathbf{t}|\mathbf{t}-1}$, the likelihood function for time t conditioned on its history through time $t - 1$ is:

$$f(\mathbf{Y}_{\mathbf{t}} | \mathcal{H}_{t-1}, \theta) = \sum_{i=1}^J \pi_{t=i|t-1} f(\mathbf{Y}_{\mathbf{t}} | S_t = i, \theta).$$

The likelihood is a function of parameters θ , the union of the set of parameters that governs the J conditional likelihood functions.

Each probability $\pi_{t=j|t-1}$ may be updated with respect to the current realization of $\mathbf{Y}_{\mathbf{t}}$ as follows:

$$(2) \quad \pi_{t=j|t} = \frac{\pi_{t=j|t-1} f(\mathbf{Y}_{\mathbf{t}} | S_t = j, \theta)}{\sum_{i=1}^J \pi_{t=i|t-1} f(\mathbf{Y}_{\mathbf{t}} | S_t = i, \theta)}.$$

Given $\xi_{\mathbf{1}|\mathbf{0}}$, recursively updating each element with (2) and projecting forward with (1) gives $\xi_{\mathbf{t}|\mathbf{t}-1}$ and $\xi_{\mathbf{t}|\mathbf{t}}$ for periods $t = 1, 2, \dots, T$, as well as the likelihood function, governed by parameters θ , the elements of P , and $\pi_{\mathbf{1}|\mathbf{0}}$.

The model would be standard if period T were the last period under consideration. However assume that there are τ additional time periods in the sample, and in these periods, $\mathbf{Y}_{\mathbf{t}}$ is not observed. A vector consisting of n substitute variables, $\mathbf{Y}_{\mathbf{t}}^f$, is observed instead, and the conditional likelihoods of these variables are governed by θ^f rather than θ . $\mathbf{Y}_{\mathbf{t}}^f$ could be a set of forecasts of $\mathbf{Y}_{\mathbf{t}}$, or a set of preliminary estimates that will be revised in a significant way at a later date.

One can interpret this model as including a set of breaks in the parameters θ , occurring

at a known point in the sample, with the transition matrix P remaining the same across the break. The log-likelihood function for this model is:

$$(3) \quad \begin{aligned} \log \mathcal{L}(\theta, \theta^f, P, \xi_{1|0}) &= \sum_{t=1}^T \log \left(\sum_{i=1}^J \pi_{t=i|t-1} f(\mathbf{Y}_t | S_t = i, \theta) \right) \\ &+ \sum_{t=T+1}^{T+\tau} \log \left(\sum_{i=1}^J \pi_{t=i|t-1} f(\mathbf{Y}_t^f | S_t = i, \theta^f) \right). \end{aligned}$$

We maintain the assumption that only the current state of the world affects the conditional densities of the \mathbf{Y}_t^f , so we may write $f(\mathbf{Y}_t^f | S_t, \theta^f)$. The usual way to proceed would be to estimate θ , θ^f , the elements of P , and $\pi_{1|0}$ by maximum likelihood of (3). Smoothed probabilities $\pi_{t=i|T+\tau}$ that use information from the entire sample could be computed using the standard algorithm of Kim (1994).

The issue in the applications considered below is small τ : there are too few observations to provide reliable estimates of θ^f through maximization of (3). In some cases of interest $\tau = 1$; with more than one state of the world, sensible parameter estimates are obviously then infeasible. However assume that we have some supplementary information at our disposal: the time series values of \mathbf{Y}_t^f for all periods $t = 1, \dots, T + \tau$. The estimation procedure outlined in the next section exploits this additional information to estimate θ^f .

3 Supplementary Estimation of the Required End-Of-Sample Parameters

The first stage of the proposed procedure side-steps the small τ problem by maximizing (3) through period T , producing parameter estimates for $\hat{\theta}$, \hat{P} and $\widehat{\pi_{1|0}}$. Hamilton (1990)

proves that the maximum likelihood estimates of θ satisfy:

$$(4) \quad \sum_{t=1}^T \sum_{i=1}^J \pi_{t=i|T} \frac{\partial \log f(\mathbf{Y}_t | S_t = i, \theta)}{\partial \theta} = 0,$$

where the derivatives are evaluated at $\hat{\theta}$, and where $\pi_{t=i|T}$ are the smoothed probabilities computed using information from the entire sample through period T . These $\pi_{t=i|T}$ are complicated functions of the model parameters, and are clearly unknown prior to estimation. However if we did treat these smoothed probabilities as given, note that the first order conditions (4) are also the first order conditions for maximization of the following log criterion function:

$$(5) \quad \log \mathcal{L}(\theta) = \sum_{t=1}^T \sum_{i=1}^J \pi_{t=i|T} \log f(\mathbf{Y}_t | S_t = i, \theta),$$

or the following criterion function:

$$(6) \quad \mathcal{L}(\theta) = \prod_{t=1}^T \prod_{i=1}^J f(\mathbf{Y}_t | S_t = i, \theta)^{\pi_{t=i|T}}.$$

With no regime-switching, so that $J = 1$, this reduces to a standard likelihood function. The smoothed probabilities provide the basis for a *within-observation* weighting of this standard likelihood, assigning some fraction $\pi_{i=t|T}$ of observation t to each state i . Since $\sum_{i=1}^J \pi_{t=i|T} = 1$ for each period t , these smoothed probabilities comprise proper weights. For some additional motivation, Appendix A draws an analogy to the use of weighted averages in another context, that of computing means from a survey sample.

These facts motivate the second stage estimation procedure for θ^f . The proposed

estimator of $\hat{\theta}^f$ minimizes:

$$(7) \quad \log \mathcal{L}(\theta^f) = \sum_{t=1}^T \sum_{i=1}^J \pi_{t=i|T} \log f(\mathbf{Y}_t^f | S_t = i, \theta^f),$$

with first order conditions solving:

$$(8) \quad \sum_{t=1}^T \sum_{i=1}^J \pi_{t=i|T} \frac{\partial \log f(\mathbf{Y}_t^f | S_t = i, \theta^f)}{\partial \theta^f} = 0.$$

As noted earlier, these $\pi_{t=i|T}$ are functions of parameters estimated in the first stage; proper but cumbersome notation would write $\pi_{t=i|T}(\hat{\theta}, \hat{P}, \widehat{\pi_{1|0}})$. Standard errors for θ^f are computed treating (7) as a standard log likelihood function, accounting for the variance of parameters estimated in the first stage entering through $\pi_{t=i|T}$ using the two-step maximum likelihood formulas in Murphy and Topel (1985).¹

With $\hat{\theta}^f$ now in hand as well as $\hat{\theta}$, \hat{P} and $\widehat{\pi_{1|0}}$ from the first stage, the usual recursion and smoothing techniques are applied to produce estimated probabilities through period $T + \tau$, with the overall likelihood for the data as in (3). Probabilities $\pi_{t=j|t-1}$ and $\pi_{t=j|t}$ through period T are produced in the standard way using (1) and (2). For periods $T + 1, \dots, T + \tau$, the updating changes to employ the likelihood of the end-of-sample data:

$$(9) \quad \pi_{t=j|t} = \frac{\pi_{t=j|t-1} f(\mathbf{Y}_t^f | S_t = j, \theta^f)}{\sum_{i=1}^J \pi_{t=i|t-1} f(\mathbf{Y}_t^f | S_t = i, \theta^f)}.$$

Smoothed probabilities $\pi_{t=i|T+\tau}$ are computed from these one-step ahead and contemporaneous probabilities using the standard algorithm of Kim (1994), with the data as

¹I use numerical derivatives to implement the Murphy and Topel (1985) formulas in the applications below.

in the likelihood (3). These techniques are identical to those that would be applied if estimation of parameters had been straightforward through maximization of (3). Only the estimation of $\widehat{\theta}^f$ is non-standard; computation of all probabilities is done in the usual way.

This is the basics of the procedure. Before moving on to the empirical applications below, a couple of additional points should be noted. First, in some situations discussed below I conducted first stage estimation on the entire sample $T + \tau$, imposing $\theta = \theta^f$ for the last τ periods. Of course, in the second stage θ^f was computed in the same manner as outlined above and so was allowed to differ from θ . This modification did not pose any problems to the basic procedure, and since the ratio $\frac{\tau}{T}$ was small in each case, the estimated $\widehat{\theta}$ in the first stage largely reflect θ rather than the θ^f that govern the last τ observations. Second, it is clear that the second-stage procedure can be repeated as many times as desired. Assume that the analyst suspects that k different types of data inhabit the last τ_k observations of the sample under consideration, so the likelihood is of the form:

$$\begin{aligned}
\log \mathcal{L} = & \sum_{t=1}^T \log \left(\sum_{i=1}^J \pi_{t=i|t-1} f(\mathbf{Y}_t | S_t = i, \theta) \right) \\
& + \sum_{t=T+1}^{T+\tau_1} \log \left(\sum_{i=1}^J \pi_{t=i|t-1} f(\mathbf{Y}_t^{\mathbf{f}_1} | S_t = i, \theta^{f_1}) \right) \\
& + \sum_{t=T+\tau_1+1}^{T+\tau_2} \log \left(\sum_{i=1}^J \pi_{t=i|t-1} f(\mathbf{Y}_t^{\mathbf{f}_2} | S_t = i, \theta^{f_2}) \right) \\
& \quad \vdots \\
(10) \quad & + \sum_{t=T+\tau_{k-1}+1}^{T+\tau_k} \log \left(\sum_{i=1}^J \pi_{t=i|t-1} f(\mathbf{Y}_t^{\mathbf{f}_k} | S_t = i, \theta^{f_k}) \right).
\end{aligned}$$

This situation could be handled easily with first stage estimation as before, and k ap-

plications of the second stage estimation solving (8), using each of the full time series of $\mathbf{Y}_t^{f_1}, \mathbf{Y}_t^{f_2} \dots \mathbf{Y}_t^{f_k}$ through period T . Probabilities are then computed as in (9) with the appropriate density, and smoothing is as before.

4 Empirical Applications

The first two applications considered below show estimation results using (8), and the third application shows estimation results and a sequence of smoothed probabilities around the start of the 2001 recession, using (9) and the algorithm of Kim (1994).

Application 1: Vintage Differences in Real GDP Growth

The U.S. Bureau of Economic Analysis (BEA) releases a number of different vintages of quarterly real GDP growth for any given quarter. The first three, the “advance”, “preliminary”, and “final” current quarterly estimates, are released about a month, two months, and three months after the end of the quarter, respectively. The “final” current quarterly estimates then stand until the summer, when the BEA releases an annual revision. Any given quarter is revised at its first three annual revisions (so 1999Q4 is revised at the 2000, 2001 and 2002 annual revisions), and then subsequently at benchmark revisions that occur approximately every five years. These revisions incorporate previously unavailable data to produce improved estimates, and the cumulative magnitude of all the revisions can be sizeable; see Mankiw and Shapiro (1986) and Fixler and Grimm (2002).

Assume that estimates that have passed through at least one annual revision are governed by a common set of parameters θ .² Conditional on the state of the world, the

²In principle we could treat differently data that have been through different numbers of annual and

distribution of real GDP growth, Δy_t , is Gaussian, so:

$$f(\Delta y_t | S_t) = \frac{1}{\sqrt{2\pi(\sigma)^2}} \exp\left(-\frac{(\Delta y_t - \mu_{S_t})^2}{2(\sigma)^2}\right).$$

There are two states of the world, expansion ($S_t = 1$) and recession ($S_t = 2$), and the mean growth rate μ_{S_t} switches depending on the state of the economy. Simplifying the notation for the estimated probabilities, label:

$$\begin{aligned} \text{prob}(S_{t-1} = 1 \mid \mathcal{H}_{t-1}) &= \pi_{t-1|t-1}, \quad \text{and:} \\ \text{prob}(S_{t-1} = 2 \mid \mathcal{H}_{t-1}) &= 1 - \pi_{t-1|t-1}. \end{aligned}$$

The Markov transition matrix contains two parameters, $p_{11} = \text{prob}(S_t = 1 | S_{t-1} = 1)$ and $p_{22} = \text{prob}(S_t = 2 | S_{t-1} = 2)$; application of the transition matrix yields:

$$\begin{pmatrix} \pi_{t|t-1} \\ 1 - \pi_{t|t-1} \end{pmatrix} = \begin{bmatrix} p_{11} & 1 - p_{22} \\ 1 - p_{11} & p_{22} \end{bmatrix} \begin{pmatrix} \pi_{t-1|t-1} \\ 1 - \pi_{t-1|t-1} \end{pmatrix}.$$

Estimating this model through period T , equation (4) yields:

$$(11) \quad \begin{aligned} \widehat{\mu}_{S_t=1} &= \frac{\sum_{t=1}^T \pi_{t|T} \Delta y_t}{\sum_{t=1}^T \pi_{t|T}} \\ \widehat{\mu}_{S_t=2} &= \frac{\sum_{t=1}^T (1 - \pi_{t|T}) \Delta y_t}{\sum_{t=1}^T (1 - \pi_{t|T})} \\ \widehat{(\sigma)^2} &= \frac{\sum_{t=1}^T \pi_{t|T} (\Delta y_t - \widehat{\mu}_{S_t=1})^2 + (1 - \pi_{t|T}) (\Delta y_t - \widehat{\mu}_{S_t=2})^2}{T}. \end{aligned}$$

These are simple weighted averages computed using the smoothed probabilities as weights.

benchmark revisions.

Consider the BEA's time series after a "final" current quarterly release. "Final" growth rates inhabit the last τ quarters of the sample - i.e. the quarters where data have not passed through an annual revision yet. The distribution of "final" current quarterly GDP growth Δy_t^f is governed by θ^f :

$$f\left(\Delta y_t^f | S_t\right) = \frac{1}{\sqrt{2\pi}(\sigma^f)^2} \exp\left(\frac{-\left(\Delta y_t^f - \mu_{S_t}^f\right)^2}{2(\sigma^f)^2}\right).$$

The likelihood is as in (3).

Consider two examples of how Δy_t^f may differ from Δy_t . First, assume each Δy_t^f is a noisy estimate of Δy_t , so:

$$\Delta y_t^f = \Delta y_t + \epsilon_t^f,$$

with ϵ_t^f orthogonal to Δy_t and with variance σ_ϵ^2 . Then $(\sigma_f)^2 = \sigma^2 + \sigma_\epsilon^2 > \sigma^2$. When updating with respect to Δy_t^f , this higher variance implies smaller probability shifts for any given growth rate; see Nalewaik (2007). For the second example, assume that Δy_t^f is not noisy, but simply misses some of the variation in mean GDP growth arising from recessions and expansions, so $\mu_{S_t=1}^f - \mu_{S_t=2}^f < \mu_{S_t=1} - \mu_{S_t=2}$. The effect is again to dampen the probability shifts when updating with respect to Δy_t^f , assuming the conditional variance of both types of estimates is the same.

After first stage estimation gives $\hat{\theta}$, \hat{P} and $\widehat{\pi_{10}}$, application of (8) yields $\hat{\theta}^f$ by taking

the same weighted averages as in (11):

$$\begin{aligned}
 (12) \quad \widehat{\mu}_{S_t=1}^f &= \frac{\sum_{t=1}^T \pi_{t|T} \Delta y_t^f}{\sum_{t=1}^T \pi_{t|T}} \\
 \widehat{\mu}_{S_t=2}^f &= \frac{\sum_{t=1}^T (1 - \pi_{t|T}) \Delta y_t^f}{\sum_{t=1}^T (1 - \pi_{t|T})} \\
 (\widehat{\sigma}^f)^2 &= \frac{\sum_{t=1}^T \pi_{t|T} \left(\Delta y_t^f - \widehat{\mu}_{S_t=1}^f \right)^2 + \sum_{t=1}^T (1 - \pi_{t|T}) \left(\Delta y_t^f - \widehat{\mu}_{S_t=2}^f \right)^2}{T}.
 \end{aligned}$$

These conditional moments $\widehat{\theta}_f$ tell us how Δy_t^f behaves during recessions and expansions, using the latest available time series to assign probabilities of recession and expansion to each date.

Having obtained each current quarterly growth rate of real GDP released by BEA since 1978Q1, my sample starts in that quarter. The two panels of table 1 show estimation results using (11) on the time series available at BEA's release of 2000Q3 and 2005Q4 data.³ Standard errors are below parameter estimates in parentheses. With a fairly short sample, the inclusion of an additional recession can swing around parameter estimates somewhat: the mean growth rate in recessions rises substantially when the relatively mild 2001 recession is included in the sample. Figure 1 shows smoothed probabilities from the second panel, the weights in (11) and (12). For the most part, these line up well with the NBER's start and end dates for recessions.⁴

The top panel of table 2 shows $\widehat{\theta}^f$ for the "final" current quarterly growth rates using

³The entire sample is used for first stage estimation here, including the last observations that have not passed through annual revisions. In the first panel, 3 out of 91 observations have not passed through annual revision, and in the second panel, 4 out of 112 observations.

⁴I classify the NBER recession dates in a somewhat expansive way; the NBER dates the beginning and end of recessions by month; if any month of a quarter is classified by the NBER as in recession, I shade the entire quarter. For example, the NBER says the last recession in the plots started in March 2001 and ended in November 2001; I shade both 2001Q1 and 2001Q4.

(12) on the sample through 2005Q4. Comparing these parameters to those in the second panel of table 1, we see no evidence that these “final” growth rates are noisy, as they have lower conditional variance than the revised growth rates. However they do pick up less of the variation in mean GDP growth arising from expansions ($\mu_{S_t=1}$) and recessions ($\mu_{S_t=2}$). The next two panels show results for the “preliminary” and “advance” growth rates. These estimates have smaller conditional variances than the “final” growth rates, and they pick up slightly less of the mean growth differences across states.

Overall it seems that treating these current quarterly growth rates as if they were heavily revised should entail little loss of accuracy in estimating probabilities of recession, as their conditional moments are not so dissimilar. However this will not necessarily be the case for other time series. For example Faust, Rogers, and Wright (2005) report that, for a number of G-7 countries, preliminary GDP growth rates contain a substantial amount of noise that is later revised away. In such cases accounting for vintage differences may be quite important in avoiding misleading inferences about whether these economies are in recession.

Application 2: Greenbook Forecasts of Real GDP Growth

This application considers adding forecasts of GDP onto the end of the BEA’s officially published time series. For example, at the end of March 1995, the BEA released its “final” current quarterly estimate of real GDP growth for 1994Q4. Plenty of data on 1995Q1 was available by the end of that March, so a forecast of 1995Q1 based on that data would likely be reasonably accurate, and a probability of recession estimated using that forecast quite informative. However, in all likelihood, the statistical properties of the forecast will differ from the statistical properties of the officially published data in prior quarters, as the officially published data are based on much more complete

information. Assuming no the vintage differences, with $\tau = 1$ for the forecast appended to the officially published time series, the likelihood is as in (3). Tacking on additional forecasts is an option as well - 1995Q2, 1995Q3, etc. The likelihood will be as in (10), with $\tau_i = 1, \forall i$, if the published time series is appended with k forecasts.

The staff of the Board of Governors of the Federal Reserve System produces its Greenbook forecasts of important economic aggregates before each of the eight meetings of the Federal Open Market Committee each year. These Greenbook forecasts have been released to the public through the end of 2000 as of the writing of this paper; this application analyzes the second Greenbook forecast produced each quarter, typically produced before FOMC meetings in March, June, September and December.⁵ The top panel of Figure 2 plots the Greenbook forecast for the current quarter, the dashed green line, with the “final” estimate of GDP growth rate for that quarter, typically released by BEA about three months later.⁶ The first panel of table 3 shows parameters estimated as in (12), with smoothed probabilities computed from the univariate GDP model using the time series available with the “final” 2000Q3 release. Comparing with the first panel of table 1, the mean growth rate in expansions is underestimated somewhat, but the mean growth rate in recessions is quite accurate, and the conditional variance is smaller than the conditional variance of GDP growth.

The second panel of Figure 2 shows one-quarter ahead Greenbook forecasts plotted again with the “final” growth rates (the one-quarter ahead forecast for 1995Q2 is reported in the March 1995 Greenbook, for example); the moments of these forecasts are shown in the second panel of Table 3. These one-quarter ahead forecasts pick up

⁵Results using the first Greenbook forecast each quarter, typically produced before FOMC meetings in January, May, August, and either October or November, are similar.

⁶Prior to 1991Q4, “final” GNP growth rates are plotted.

a considerable amount of business cycle variation in real GDP growth. The two- and three-quarter ahead forecasts are plotted in Figure 3, with estimated moments in the last two panels of Table 3. As expected, the gap between the recession and expansion mean growth rates diminishes at longer forecast horizons, as it becomes more difficult to forecast the state of the economy farther into the future. As this gap between the mean growth rates converges to zero, the updating (9) converges to no updating at all, as the forecasts are not informative about the future state of the world. To the extent that these forecasts are informative, though, the accuracy of the estimated probability of recession for the *current* quarter will be improved by smoothing using the information in forecasts for future quarters. So even if the goal is simply to estimate the probability of recession for the current quarter as accurately as possible, as opposed to estimating probabilities of recession for future quarters, forecasts may be informative. However it is clear from these tables that when appending published time series with forecasts several quarters ahead, the differential information content of the forecasts vis-a-vis official data needs to be taken into account when computing probabilities.

Before continuing, it should be mentioned that parameter stability is an obvious concern here: if the techniques used to produce Greenbook forecasts have changed over time, parameters estimated from historical forecasts may not be useful in evaluating the information content in current forecasts. This is probably more of a concern for the one- to three-quarter ahead forecasts than for forecasts of the current quarter, which are mostly a matter of carefully adding up publicly-released data that BEA uses to compile its estimates. This adding up is unlikely to have changed much over time.

Application 3: Vintage Differences and Forecasts in a Bivariate Model

Nalewaik (2007) uses real time data to show that since 1978, the growth rate of gross

domestic income (GDI), deflated by the GDP deflator, has done a better job recognizing recessions than has the growth rate of real GDP. That paper recommends a bivariate Markov switching model employing both GDP and GDI; this next application studies such a model, but goes further in accounting for vintage differences in the current quarterly estimates as in application 1, and appending each BEA time series with forecasts as in application 2. Probabilities of recession are computed at different BEA release dates, with the appended forecast employing one of three different information sets depending on whether the release is an “advance,” “preliminary,” or “final”; the information set contains variables typically available at the time of that type of release. The forecasts are for the current quarter only: for example at the time of BEA’s “advance” current quarterly release of 1994Q4 data - January 1995 - the published time series is appended with a forecast for 1995Q1. These forecasts employ real time data on industrial production, employment, initial claims for unemployment insurance, and the Federal Reserve Bank of Philadelphia’s Business Outlook Survey; details of forecast computation are in Appendix B. The predicted values from the forecasts, made at the time of the “advance,” “preliminary,” and “final” releases of the previous quarter’s data, respectively, are plotted with “final” vintage data in Figures 4 to 6. These forecasts track the business cycle variation in real GDP and real GDI growth reasonably well.

A more thorough outline of the bivariate model can be found in Nalewaik (2007). Briefly, real GDP growth Δy_t^1 and real GDI growth Δy_t^2 are jointly normally distributed as:

$$\begin{bmatrix} \Delta y_t^1 \\ \Delta y_t^2 \end{bmatrix} \sim N \left(\begin{bmatrix} \mu_{S_t}^1 \\ \mu_{S_t}^2 \end{bmatrix}, \begin{bmatrix} (\sigma^1)^2 & \sigma^{12} \\ \sigma^{12} & (\sigma^2)^2 \end{bmatrix} \right).$$

The Markov transition matrix is as in application 1. Applying (4), the maximum

likelihood estimates of the parameters of the model satisfy:

$$\begin{aligned}
(13) \quad \widehat{\mu}_{S_t=1}^i &= \frac{\sum_{t=1}^T \pi_{t|T} \Delta y_t^i}{\sum_{t=1}^T \pi_{t|T}}, & \text{for } i = 1, 2; \\
\widehat{\mu}_{S_t=2}^i &= \frac{\sum_{t=1}^T (1-\pi_{t|T}) \Delta y_t^i}{\sum_{t=1}^T (1-\pi_{t|T})}, & \text{for } i = 1, 2; \\
(\widehat{\sigma}^i)^2 &= \frac{\sum_{t=1}^T \pi_{t|T} (\Delta y_t^i - \widehat{\mu}_{S_t=1}^i)^2 + (1-\pi_{t|T}) (\Delta y_t^i - \widehat{\mu}_{S_t=2}^i)^2}{T}, & \text{for } i = 1, 2; \\
\widehat{\sigma}^{12} &= \frac{\sum_{t=1}^T \pi_{t|T} (\Delta y_t^1 - \widehat{\mu}_{S_t=1}^1) (\Delta y_t^2 - \widehat{\mu}_{S_t=1}^2) + (1-\pi_{t|T}) (\Delta y_t^1 - \widehat{\mu}_{S_t=2}^1) (\Delta y_t^2 - \widehat{\mu}_{S_t=2}^2)}{T}.
\end{aligned}$$

As before, first stage maximum likelihood estimation of the model gives $\widehat{\theta}$, \widehat{P} and $\widehat{\pi}_{1|0}$. Applying (8) to each different vintage type or forecast at the end of the sample yields $\widehat{\theta}^{f_k}$; again these are simple weighted averages, using the same weights as in (13):

$$\begin{aligned}
(14) \quad \widehat{\mu}_{S_t=1}^{i,f_k} &= \frac{\sum_{t=1}^T \pi_{t|T} \Delta y_t^{i,f_k}}{\sum_{t=1}^T \pi_{t|T}}, & \text{for } i = 1, 2; \\
\widehat{\mu}_{S_t=2}^{i,f_k} &= \frac{\sum_{t=1}^T (1-\pi_{t|T}) \Delta y_t^{i,f_k}}{\sum_{t=1}^T (1-\pi_{t|T})}, & \text{for } i = 1, 2; \\
(\widehat{\sigma}^{i,f_k})^2 &= \frac{\sum_{t=1}^T \pi_{t|T} (\Delta y_t^{i,f_k} - \widehat{\mu}_{S_t=1}^{i,f_k})^2 + (1-\pi_{t|T}) (\Delta y_t^{i,f_k} - \widehat{\mu}_{S_t=2}^{i,f_k})^2}{T}, & \text{for } i = 1, 2; \\
\widehat{\sigma}^{12,f_k} &= \frac{\sum_{t=1}^T \pi_{t|T} (\Delta y_t^{1,f_k} - \widehat{\mu}_{S_t=1}^{1,f_k}) (\Delta y_t^{2,f_k} - \widehat{\mu}_{S_t=1}^{2,f_k}) + (1-\pi_{t|T}) (\Delta y_t^{1,f_k} - \widehat{\mu}_{S_t=2}^{1,f_k}) (\Delta y_t^{2,f_k} - \widehat{\mu}_{S_t=2}^{2,f_k})}{T}.
\end{aligned}$$

The top panel of Table 4 shows estimation results on the latest available time series as of the BEA's release of 2005Q4 data. The inclusion of real GDI growth in the Markov switching model has interesting effects discussed in more detail in Nalewaik (2007): the mean duration of recessions increases, leading to greater, actually positive, estimated mean growth rates in recessions. Figure 7 plots smoothed probabilities of recession from this model, and the more expansive definition of recessions is evident: the 1980 recession includes most of 1979, the 1990-1991 recession includes most of 1989, and the

2001 recession includes part of 2000 and all of 2002. With these smoothed probabilities as weights, real GDI has a smaller conditional variance than real GDP, and the gap between its high- and low-growth means is larger, implying that GDI is a better indicator of the state of the economy in this model.

The second panel of Table 4 shows $\hat{\theta}^f$ for the “final” current quarterly growth rates computed from (14).⁷ Compared to the first panel, the parameters for real GDI growth are closer to those for real GDP growth: while the conditional variance remains smaller and the gap between the high- and low-growth means remains larger, these advantages of real GDI as a signal of the state of the world are diminished. This is an important caveat to the results in Nalewaik (2007) showing that GDI outperforms GDP in recognizing the start of recessions.⁸

The third and fourth panels of Table 4 show $\hat{\theta}^f$ for the “preliminary” current quarterly growth rates. In the BEA’s “preliminary” estimates for fourth quarters, corporate profits and a couple other components of GDI are left missing because of poor data quality. These gaps were filled in with predictions as described in Appendix B. In the third panel, the time series of actual BEA releases is employed in (14), with predicted values substituting for actual releases only in fourth quarters; these parameters are similar to those for the “final” current quarterly growth rates. The fourth panel shows results through 2005Q4 using the full time series of “preliminary” predicted values. The missing information on corporate profits is important, as the gap between the high- and low-

⁷For notational simplicity, these tables write $\hat{\theta}^f$ instead of $\hat{\theta}^{f_k}$, but the likelihood is as in (10) in this section, with multiple different types of end-of-sample data.

⁸Fixler and Nalewaik (2005) studied various models characterizing the variance of GDP growth and GDI growth, and found that if what they call the pure “news” model is true, then the information content of GDI relative to GDP is lower in the early-vintage estimates than in the later vintage estimates that have passed through annual and benchmark revisions. The result here, showing that GDI has less of an advantage over GDP in identifying recessions in the early-vintage estimates, is consistent with that result from the pure “news” model.

growth means for GDI growth contracts appreciably, and is now smaller than the gap for GDP.⁹ The last panel here shows $\hat{\theta}^f$ for the “advance” current quarterly growth rates; these employ predicted values as they have the same missing components as fourth quarter “preliminary” releases, so it is not surprising that these parameter estimates are similar to those in the prior panel.

Table 5 shows estimates of $\hat{\theta}_f$ for the forecasts. For each of the three information sets we employ, there is some spread between the high- and low-growth means, confirming the visual evidence in Figures 4-6 that the estimates track the business cycle variation in the data. These mean spreads are considerably smaller than those of the official BEA estimates, and the conditional variances are about half those reported in the top panel of table 4; a differential statistical treatment of the forecasts when computing probabilities of recession seems appropriate in this case.

Figures 8-10, show smoothed probabilities of recession at selected BEA data releases around the start of the 2001 recession, computed using real time data and the paper’s new estimation techniques. Each graph shows two time series of smoothed probabilities using data from an “advance” or a “final” release for a quarter, starting with the “advance” 2000Q4 release in the top panel of Figure 8. The solid line, in black, shows smoothed probabilities that account for vintage differences but employs no forecasts of GDP or GDI or its components. In “advance” releases where GDI is unavailable, the probability of recession for that quarter is updated with respect to GDP alone using the two-stage updating procedure described in Nalewaik (2007). The dashed line in

⁹In the real time estimates of the model shown in Figures 8-12, when a predicted value is the last time series observation in the sample - i.e. when the sample period ends in a fourth quarter, a full time series of predicted values is employed to compute parameters in equation (14); when the sample ends in a first, second, or third quarter, so the last time series observation in the sample is the actual value from a BEA “preliminary” release, the time series of actual BEA releases is employed in (14), with predicted values substituting for actual releases only in fourth quarters.

blue shows smoothed probabilities that account for vintage differences and employ the forecasts as well, extending the time series forward one additional quarter.¹⁰ To get a better sense of the timing involved here, written in the label for each time series is the month when that data release occurred.

The “advance” reading for 2000Q4 real GDP growth released in January 2001 was somewhat weak at 1.4%, leading to a probability of recession of about 10% in the basic model with vintage differences but no forecasts. However the real time forecasts for 2001Q1 real GDP and GDI growth were weaker, predicting around-zero growth, as initial claims for unemployment insurance had already started to shoot up and the Business Outlook Survey was extraordinarily weak. The model incorporating these forecasts shows estimated probabilities of recession of 43% in 2001Q1 and 22% in 2000Q4.

Little changes with the BEA’s “preliminary” 2000Q4 release in February. The BEA’s March 2001 “final” release brings the first published value of real GDI growth in 2000Q4, an anemic reading of 0.3% growth. The bottom panel of Figure 8 shows smoothed probabilities for this release: the weak value for GDI growth brings the 2000Q4 recession probability computed without forecasts up to about the same level as the probability computed with the aid of the forecasts.

Moving to Figure 9, the April 2001 release of “advance” 2001Q1 real GDP growth, 2.0%, adds little information about whether or not the economy is in recession. The 2001Q2 forecasts for GDP and GDI growth of 1.2% and 1.4% are weaker, though again they do not give an unambiguous signal of recession: the 2001Q2 probability of recession from the model with forecasts is 54%, and the smoothed 2001Q1 probability is 45%,

¹⁰In the likelihood (10) for these estimates, $k = 2$ for the “final” current quarterly releases, and $k = 3$ for the other two types of releases, since some quarters of BEA data before the last quarter will be “final” vintage, the last quarter will be “advance” or “preliminary”, and appended to the published time series is the forecast for the current quarter.

compared to the 27% probability of recession computed without the forecasts.

The May 2001 “preliminary” release brings the BEA’s first estimate of 2001Q1 real GDI growth, and its 0.1% value leads all the recession probabilities to jump up dramatically. Smoothed probabilities from the June “final” release are similar; these are plotted in the bottom panel of Figure 9. Although these May and June releases bring the first clear evidence of recession whether or not the forecasts are employed, the forecasts allow the model to give a clearer warning signal from January to April. Furthermore, Figure 10 shows probabilities of recession from the September “final” release of 2001Q2 data; it is notable that the probability of recession ticks down if we do not employ the forecasts, perhaps leading to some additional doubt about whether or not the economy is in recession. If we employ the forecasts, however, there is little doubt, with estimated probabilities of recession in the first three quarters of the year consistently above 80%.¹¹

These results indicate that even a simple, small-scale forecasting model for GDP and GDI growth one quarter ahead has the potential to significantly increase the timeliness of estimated probabilities of recession from basic Markov switching models. More comprehensive forecasting models could presumably do an even better job.

5 Alternative Procedures: Discussion

The multi-stage nature of the procedure outlined in this paper may be unappealing to some readers, and since the procedure relies on the accessibility of the full substitute time series \mathbf{Y}_t^f for periods $t = 1, \dots, T + \tau$, an alternative would be to simply model the joint distribution of the time series of interest \mathbf{Y}_t with \mathbf{Y}_t^f . Conditional on the state of the

¹¹The release of negative GDP and GDI growth in 2001Q3 makes the recession clear in both sets of probabilities.

world, the distribution is $f(\mathbf{Y}_t, \mathbf{Y}_t^f | S_t = i, \theta, \theta^f)$. If \mathbf{Y}_t^f is a forecast, for example a linear function of other variables as in application 3 (so $\mathbf{Y}_t^f = \mathbf{X}_t \beta$), use of a joint distribution such as $f(\mathbf{Y}_t, \mathbf{X}_t | S_t = i, \theta, \theta^x)$ is an option as well. As before, \mathbf{Y}_t is unobserved for the last τ periods of the sample, but this poses no problems: the marginal distributions of \mathbf{Y}_t^f or \mathbf{X}_t can be employed to partially update probabilities in those last τ periods, with the remainder of the updating occurring later when \mathbf{Y}_t is observed. An outline of this procedure is in Nalewaik (2007).

When dealing with different vintages of estimates as in application 1, this option is clearly unappealing. After a revision to a quarterly estimate of GDP or GDI, the BEA essentially throws out the prior estimate; since the new revised estimate incorporates more information and is more accurate, there is no reason to continue to examine the prior one. A model such as $f(\mathbf{Y}_t, \mathbf{Y}_t^f | S_t = i, \theta)$, then, would allow the probabilities of recession to depend on inferior data. For forecasts as in application 2, the situation is analogous: once the variable to be forecast becomes known, the prior forecast typically becomes uninteresting. The forecast is clearly less accurate than the realization, so a model such as $f(\mathbf{Y}_t, \mathbf{Y}_t^f | S_t = i, \theta, \theta^f)$ would again allow the probabilities of recession to depend on inferior data. When it is feasible to decompose the forecast into its component indicator variables \mathbf{X}_t as in application 3, this remains true. GDP and GDI are the most comprehensive measures of the state of the economy. A model such as $f(\mathbf{Y}_t, \mathbf{X}_t | S_t = i, \theta, \theta^x)$ places the indicators \mathbf{X}_t on an equal footing with GDP and GDI, giving them undue importance. For example, the results in application 3 show that the Federal Reserve Bank of Philadelphia's Business Outlook Survey, a monthly diffusion index computed from a survey of mid-Atlantic-based manufacturers, is somewhat useful as an early indicator of the state of the economy before much other data on the period is available. Using such a survey for forecasting in that way seems reasonable, but

placing it on an equal footing with GDP or GDI in computing historical estimates of probabilities of recession is an entirely different matter.

The multi-step nature of the procedure outlined in this paper has some appeal as well, as it allows the econometrician to handle an arbitrary number of different types of end-of-sample data at essentially zero computational cost. Often nothing more complicated than taking weighted means and variances is required to compute parameters estimated in stages beyond the first, although computation of standard errors is more complicated. Furthermore, the sequential estimation procedure may reduce parameter estimation error compared to such a case where all parameters are estimated jointly. Consider the time series after an “advance” or “preliminary” release in application 3, described as a bivariate model with three different types of data at the end of the sample. A joint distribution of these data, conditional on the state of the world, is:

$$f(\mathbf{Y}_t, \mathbf{Y}_t^{f_1}, \mathbf{Y}_t^{f_2}, \mathbf{Y}_t^{f_3} | S_t = i, \theta, \theta^{f_1}, \theta^{f_2}, \theta^{f_3}).$$

Including all the cross covariance parameters between the different types of data, this model contains 55 parameters; simultaneous estimation of all these model parameters would likely lead to considerable parameter estimation error. Under the sequential estimation employed in application 3, no more than 10 parameters are estimated at any stage, and estimation of 24 of the cross covariance parameters is unnecessary.

While modelling the joint distribution $f(\mathbf{Y}_t, \mathbf{Y}_t^f | S_t = i, \theta, \theta^f)$ is the most obvious alternative to the procedures outlined in this paper, some more complicated models have been considered as well in attempting to incorporate supplementary information into Markov switching models. For example, Diebold, Lee and Weinbach (1994), Filardo (1994), Filardo and Gordon (1998) and others have modelled the parameters in

the transition matrix P as functions of variables such as the \mathbf{X}_t in application 3, and Hamilton and Perez-Quiros (1996) have studied models where the state of the world shifts at different times for different variables. The concerns discussed earlier, of estimated probabilities depending on inferior data, are less of a concern in these models, as the most comprehensive variables like GDP and GDI can be treated differently from the variables used to forecast them. However some concerns along these lines remain. In addition, the added complexity of these models may reduce their robustness; for example Hamilton (2005) reports that models where the transition probabilities are functions of variables suffer from imprecise parameter estimates. A virtue of the procedures discussed in sections 2-4 is their simplicity; as the different applications show, the simplicity and robustness of the estimation techniques allows them to handle complicated situations with relative ease.

6 Conclusions

This paper outlines techniques to account for differences between the statistical properties of different time series observations, in the context of the class of Markov switching models discussed in Hamilton (1994). It focuses on the case where a small number of observations at the end of the time series of interest exhibit statistical properties that are different from the rest, with this number being too small for standard maximum likelihood techniques to produce precise parameter estimates. The supplementary estimation techniques proposed here are computationally simple and highly intuitive: they are often simple generalizations of taking means and variances in the model's different states of the world.

Applications focused on a Markov switching model where the U.S. economy, as mea-

sured by GDP and its income-side counterpart GDI, switches between recession and expansion. One application shows how to treat differences across vintages of GDP, allowing early-vintage estimates that have not passed through many revisions to have different statistical properties than later-vintage estimates that have passed through a number of revisions. Perhaps more interesting, the techniques outlined here allow estimated probabilities of recession to reflect the information about the state of the economy embedded in forecasts. To do so, the forecasts are simply appended to the end of officially published time series, and the sample period is extended to include those additional periods. The statistical properties of the forecasts are, of course, allowed to differ in an arbitrary way from those of the official data, but to the extent that the forecasts are informative about the state of the economy in the forecast period, their incorporation will improve the accuracy of estimated probabilities of recession. Empirical results bear this out, as an application studying a bivariate model of GDP and GDI shows that the incorporation of forecasts improved the timeliness of the model in signalling that the economy was entering recession in early 2001.

Table 1:
Real Time Estimates of Markov Switching Model for GDP
1978Q1-2000Q3

$\mu_{S=1}$	$\mu_{S=2}$	$(\sigma)^2$	p_{11}	p_{22}	$p_{1,start}$
4.05	-1.41	7.60	0.95	0.73	1.00
(0.38)	(1.52)	(1.23)	(0.03)	(0.18)	(1.55)

1978Q1-2005Q4

$\mu_{S=1}$	$\mu_{S=2}$	$(\sigma)^2$	p_{11}	p_{22}	$p_{1,start}$
3.91	-0.40	7.04	0.95	0.76	1.00
(0.45)	(1.47)	(1.06)	(0.03)	(0.16)	(1.55)

Table 2:
Parameters for Current Quarterly Estimates, 1978Q1-2005Q4,
Evaluated Using Probabilities Computed from Latest Available Data

$\mu_{S=1}^f$	$\mu_{S=2}^f$	$(\sigma^f)^2$
Final		
3.48	-0.15	6.06
(0.39)	(1.43)	(0.88)
Preliminary		
3.46	-0.23	5.67
(0.39)	(1.48)	(0.84)
Advance		
3.27	-0.22	5.24
(0.37)	(1.43)	(0.77)

Table 3:
Parameters for Greenbook Predictions of Real GDP Growth, 1978Q1-2000Q3,
 Evaluated Using Probabilities Computed from Latest Available Data

Current Quarter, 2nd Greenbook in Quarter

$\mu_{S=1}^f$	$\mu_{S=2}^f$	$(\sigma^f)^2$
3.02	-1.53	4.53
(0.30)	(1.40)	(0.74)

One Quarter Ahead, 2nd Greenbook in Quarter

$\mu_{S=1}^f$	$\mu_{S=2}^f$	$(\sigma^f)^2$
2.61	-0.32	3.59
(0.29)	(0.77)	(0.53)

Two Quarters Ahead, 2nd Greenbook in Quarter

$\mu_{S=1}^f$	$\mu_{S=2}^f$	$(\sigma^f)^2$
2.59	0.73	2.33
(0.21)	(0.65)	(0.36)

Three Quarters Ahead, 2nd Greenbook in Quarter

$\mu_{S=1}^f$	$\mu_{S=2}^f$	$(\sigma^f)^2$
2.60	1.32	1.72
(0.17)	(0.53)	(0.27)

Table 4:
Real Time Estimates of Markov Switching Model for GDP and GDI, 1978Q1-2005Q4

Parameters for Latest Available Data

$\mu_{S=1}^1$	$\mu_{S=1}^2$	$\mu_{S=2}^1$	$\mu_{S=2}^2$	$(\sigma^1)^2$	$(\sigma^2)^2$	σ^{12}	p_{11}	p_{22}	$p_{1,start}$
4.18	4.36	0.63	0.26	7.20	6.03	4.96	0.93	0.85	1.00
(0.34)	(0.34)	(0.67)	(0.64)	(1.03)	(0.90)	(0.85)	(0.03)	(0.08)	(1.28)

Parameters for Current Quarterly Estimates

Evaluated Using Probabilities Computed from Latest Available Data

$\mu_{S=1}^{1,f}$	$\mu_{S=1}^{2,f}$	$\mu_{S=2}^{1,f}$	$\mu_{S=2}^{2,f}$	$(\sigma^{1,f})^2$	$(\sigma^{2,f})^2$	$\sigma^{12,f}$
Final						
3.72	3.96	0.69	0.67	6.11	5.71	5.52
(0.30)	(0.30)	(0.69)	(0.68)	(0.82)	(0.79)	(0.78)

Preliminary,

GDI predicted values used in 4th quarters only

3.71	3.92	0.67	0.65	5.84	5.47	5.29
(0.29)	(0.30)	(0.69)	(0.66)	(0.78)	(0.77)	(0.75)

Preliminary,

GDI Predicted values used in all quarters

3.68	3.67	0.68	1.07	5.82	5.13	4.20
(0.29)	(0.28)	(0.72)	(0.60)	(0.77)	(0.68)	(0.63)

Advance,

GDI Predicted values used in all quarters

3.49	3.67	0.65	0.99	5.35	5.12	4.21
(0.29)	(0.29)	(0.70)	(0.62)	(0.71)	(0.68)	(0.62)

Note: Moments of GDP have a 1 superscript, while moments of GDI have a 2 superscript.

Table 5:
Parameters for Predictions of GDP and GDI One Quarter Ahead, 1978Q1-2005Q4,
 Evaluated Using Probabilities Computed from Latest Available Data

Predicted at time of previous Final Current Quarterly Release
 (Explanatory Variables: Industrial Production Growth, Employment Growth,
 Philadelphia Fed Business Outlook Survey)

$\mu_{S=1}^{1,f}$	$\mu_{S=1}^{2,f}$	$\mu_{S=2}^{1,f}$	$\mu_{S=2}^{2,f}$	$(\sigma^{1,f})^2$	$(\sigma^{2,f})^2$	$\sigma^{12,f}$
3.41	3.51	0.47	0.45	3.36	3.75	3.54
(0.23)	(0.25)	(0.58)	(0.61)	(0.44)	(0.48)	(0.46)

Predicted at time of previous Preliminary Current Quarterly Release
 (Explanatory Variables: Industrial Production Growth, New Claims for Unemployment Insurance,
 Philadelphia Fed Business Outlook Survey)

$\mu_{S=1}^{1,f}$	$\mu_{S=1}^{2,f}$	$\mu_{S=2}^{1,f}$	$\mu_{S=2}^{2,f}$	$(\sigma^{1,f})^2$	$(\sigma^{2,f})^2$	$\sigma^{12,f}$
3.46	3.63	0.90	0.84	2.84	3.23	3.02
(0.21)	(0.23)	(0.56)	(0.60)	(0.24)	(0.25)	(0.24)

Predicted at time of previous Advance Current Quarterly Release
 (Explanatory Variables: New Claims for Unemployment Insurance,
 Philadelphia Fed Business Outlook Survey)

$\mu_{S=1}^{1,f}$	$\mu_{S=1}^{2,f}$	$\mu_{S=2}^{1,f}$	$\mu_{S=2}^{2,f}$	$(\sigma^{1,f})^2$	$(\sigma^{2,f})^2$	$\sigma^{12,f}$
3.33	3.69	0.70	0.87	2.93	3.29	3.10
(0.22)	(0.23)	(0.54)	(0.57)	(0.20)	(0.21)	(0.20)

Note: Moments of GDP have a 1 superscript, while moments of GDI have a 2 superscript.

Appendix A: Comparison with Simple Weighted Averages

To better understand equation (6), perhaps it would be useful to consider for a moment a problem quite different from the one at hand, one where the use of weighted averages is widespread. The problem is that of estimating the population mean and variance of a variable X from a sample $i = 1, 2, \dots, I$ of population members X_i , and where individuals' probabilities of selection into the sample differ. This problem is ubiquitous in the statistics literature on sampling; if the probabilities of selection are known, weights π_i for each population member may be constructed, and the population mean and variance are estimated straightforwardly as:

$$(A.1) \quad E(X) = \hat{\mu} = \frac{\sum_{i=1}^I \pi_i X_i}{\sum_{i=1}^I \pi_i} \quad \text{and:}$$

$$(A.2) \quad \text{var}(X) = \hat{\sigma}^2 = \frac{\sum_{i=1}^I \pi_i (X_i - \hat{\mu})^2}{\sum_{i=1}^I \pi_i}.$$

Assume that X follows a normal distribution with parameters μ and σ^2 , so:

$$f(X_i) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(X_i - \mu)^2}{2\sigma^2}\right).$$

Given this distributional assumption, we can work backward and ask what likelihood function produces formulas (A.2) and (A.3); the answer is:

$$(A.3) \quad \mathcal{L}(X_1, X_2, \dots, X_I, \mu, \sigma^2) = \prod_{i=1}^I f(X_i)^{\pi_i}.$$

This is a standard likelihood function except for the weights π_i ; we could employ some type of normalizing constant to account for the presence of the weights in the likelihood, but this would impact neither the maximum likelihood estimates nor their standard

errors. Since the likelihood is a multiplicative function of the $f(X_i)$, the π_i enter the likelihood as geometric weights. This is intuitive; as an example consider the portion of the likelihood for observations i and $i + 1$, and say $\pi_i = 1$ and $\pi_{i+1} = 2$, so:

$$\begin{aligned} \mathcal{L}(X_i, X_{i+1}, \mu, \sigma^2) &= f(X_i)^{\pi_i} f(X_{i+1})^{\pi_{i+1}} \\ &= f(X_i) f(X_{i+1})^2 \\ &= f(X_i) f(X_{i+1}) f(X_{i+1}). \end{aligned}$$

The weights indicate that observations of type $i + 1$ are twice as prevalent in the overall population as observations of type i , so as the above manipulations indicate, the weights effectively produce an extra observation for X_{i+1} . The general logic carries over to the case where the weights take on any positive values, not just integers.

Returning to the problem at hand and the likelihood (6), the similarity to (A.3) is evident. The major difference is that the Markov switching model implements a *within-observation* weighting; for the two-state case, the model assigns some fraction $\pi_{t|T}$ of observation t to the expansion state, and the remaining $1 - \pi_{t|T}$ of observation t to the recession state.

Appendix B: Computation of Forecasts for Real GDP and Real GDI Growth

Some components of GDI are missing in some BEA releases; before proceeding to the forecasts, this Appendix discusses computation of the predicted values for these missing components. Corporate profits is missing in the “advance” releases and fourth-quarter “preliminary” releases; to predict these, a time series consisting of the growth rates of “final” vintage corporate profits is regressed on the appropriate vintage (either

“advance” or “preliminary” depending on which set of missing values we are trying to fill) of other available components of GDP and GDI.¹² These explanatory variables are the growth rates of private consumption of fixed capital, and a couple of other cyclically sensitive variables - residential investment and personal consumption expenditures for durable goods.¹³ Since 1992, net factor income from abroad has been missing in the same releases as corporate profits, and from 1992 to 2002 net interest is missing in these releases as well. Predicted values for these two variables are computed in the same way in a similar fashion; for net factor income the explanatory variables are taxes on production and imports less subsidies and two lags of net factor income (“final” vintage), and for net interest the explanatory variables are one of its lags and employee compensation growth. These regressions are obviously somewhat ad hoc, and are just meant to be a rough attempt to fill in the missing data.

The forecasts are then constructed as follows. At the time of the “advance” release for quarter t , almost a months worth of weekly data on new claims for unemployment insurance in quarter $t + 1$ is available. The Federal Reserve Bank of Philadelphia’s Business Outlook Survey for the first month of the quarter is typically available as well; the explanatory variables in the “advance” information set are the Survey’s diffusion index measuring general business activity and the average value for claims.¹⁴ At the time of the “preliminary” release for quarter t , initial claims and the Business Outlook Survey are available for the first two months of the quarter. In addition, a wide range

¹²The use of the full time series of early release estimates as a dependent variable follows Koenig, Dolmas, and Piger (2003).

¹³The series of predicted values consists of in-sample predicted values from 1978 to 1988, and out-of-sample predicted values from rolling regressions from 1989 to 2005.

¹⁴Both the claims series and the Philadelphia Fed business survey are little revised except for updates to seasonal factors; I seasonally adjust these series on a rolling basis myself to ensure that only information in time t is employed in computing seasonal factors for time series observation t .

data on the first month of the quarter is available, including industrial production and non-farm payroll employment;¹⁵ employment growth added little explanatory power, so I used the average of the two months of initial claims, the latest available Business Outlook Survey, and the latest available three month growth rate of industrial production as explanatory variables. One month of additional data on all these variables is available at the time of the “final” release for quarter t ; the information set employed there is the latest three month growth rates of industrial production and payroll employment (growth rates from the middle month of quarter t to the middle month of quarter $t + 1$), and the latest available Business Outlook Survey, from the last month of quarter $t + 1$. A history of predicted values was stored from rolling regressions forecasting GDP and GDI using each information set.¹⁶

Although the in-sample R^2 values from the forecasting regressions over the full sample were not particularly high, ranging from 0.34 to 0.63, Figures 3 to 5 show that the predicted values track the declines in GDP and GDI growth during the recessions in the sample. In addition, with the exception of an episode in mid-1995, these predicted values do not turn down in expansions.

¹⁵I employ real-time data on these variables from the Federal Reserve Bank of Philadelphia’s dataset.

¹⁶As before, the time series of predicted values consists of in-sample predicted values from 1978 to 1988, and out-of-sample predicted values from the rolling regressions from 1989 to 2005. For GDP, the time series of either “advance”, “preliminary”, or “final” GDP growth is regressed on the appropriate set of variables (again the strategy suggested in Koenig, Dolmas, and Piger (2003)). For GDI, “final” vintage growth rates are the dependent variable in all specifications; since “final” GDI for period t is not available with the “advance” and “preliminary” time t releases, the rolling regressions employ an in-sample period through $t - 1$ only.

References

- [1] Diebold, Francis X.; Lee, Joon-Haeng; and Weinbach, Gretchen C. “Regime Switching with Time-Varying Transition Probabilities,” in *Nonstationary Time Series Analysis and Cointegration*, edited by C. Hargreaves, Oxford, 1994.
- [2] Filardo, Andrew J. “Business-Cycle Phases and Their Transitional Dynamics” *Journal of Business and Economic Statistics*, 1994, 299-308.
- [3] Filardo, Andrew J. and Gordon, Stephen F. “Business Cycle Duration” *Journal of Econometrics*, 1998 (85), pp. 99-123.
- [4] Fixler, Dennis J. and Grimm, Bruce T. “Reliability of GDP and Related NIPA Estimates” *Survey of Current Business*, January 2002, pp. 9-27.
- [5] Fixler, Dennis J. and Nalewaik, Jeremy J. “News, Noise, and Estimates of the True Unobserved State of the Economy”, BEA working paper 2004-08, 2004.
- [6] Hamilton, James D. “Analysis of Time Series Subject to Changes in Regime” *Journal of Econometrics*, 1990 (45), pp. 39-70.
- [7] Hamilton, James D. *Time Series Analysis*. Princeton, Princeton University Press, 1994.
- [8] Hamilton, James D. “Regime-Switching Models.” Prepared for: *Palgrave Dictionary of Economics*, 2005.
- [9] Hamilton, James D. and Perez-Quiros, Gabriel. “What Do the Leading Indicators Lead?” *Journal of Business*, 1996 (69), pp. 27-49.

- [10] Kim, Chang-Jin. “Dynamic Linear Models with Markov Switching” *Journal of Econometrics*, 1994 (60), pp. 1-22.
- [11] Koenig, Evan; Dolmas, Sheila; and Piger, Jeremy. “The Use and Abuse of Real-Time Data in Economic Forecasting.” *Review of Economics and Statistics*, 2003 (85), pp. 618-628.
- [12] Murphy, Kevin, and Topel, Robert. “Estimation and Inference in Two-Step Econometric Models;” *Journal of Business and Economic Statistics*, 1985, 370-379.
- [13] Nalewaik, Jeremy J. “Estimating Probabilities of Recession in Real Time Using GDP and GDI.” FEDS working paper 2007-07, 2007.

Figure 1: Smoothed Probabilities of Recession




Figure 2: Greenbook Forecasts and Actual "Final" Current Quarterly GDP Growth





Figure 3: Greenbook Forecasts and Actual "Final" Current Quarterly GDP Growth

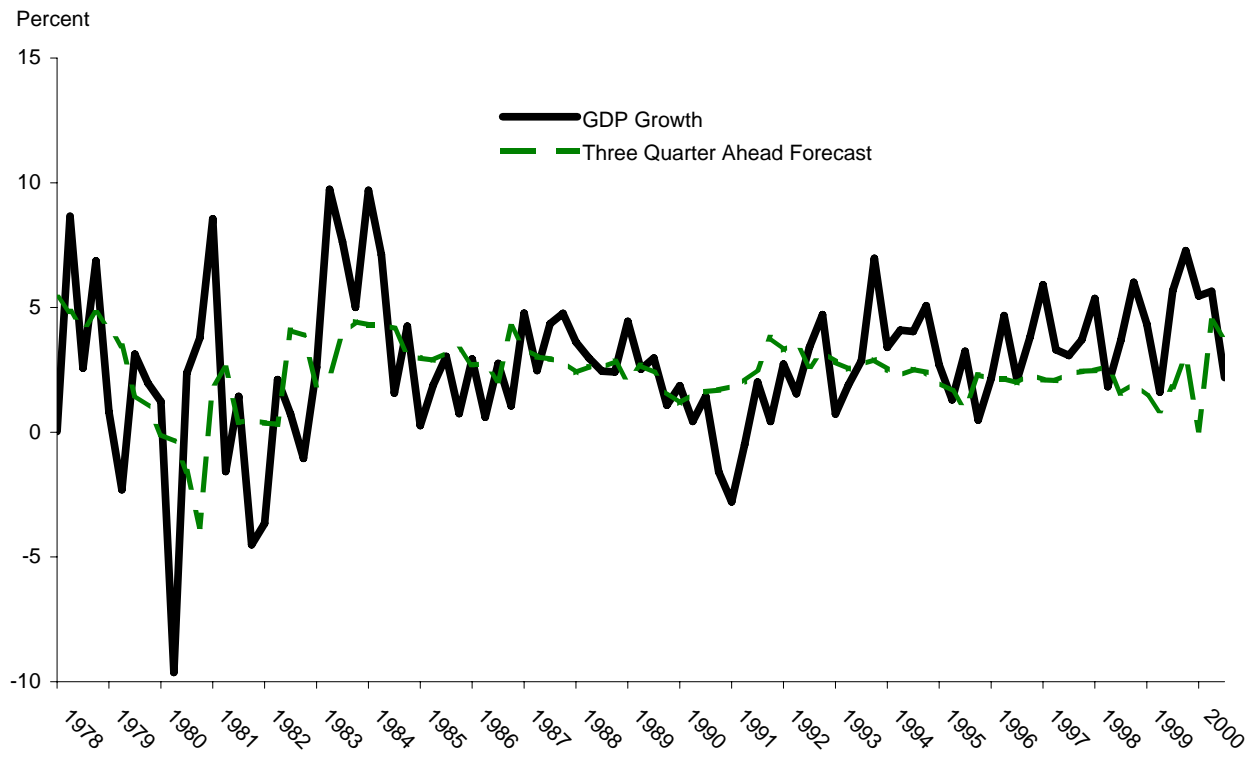



Figure 4: Predicted and Actual Real "Final" Current Quarterly Data


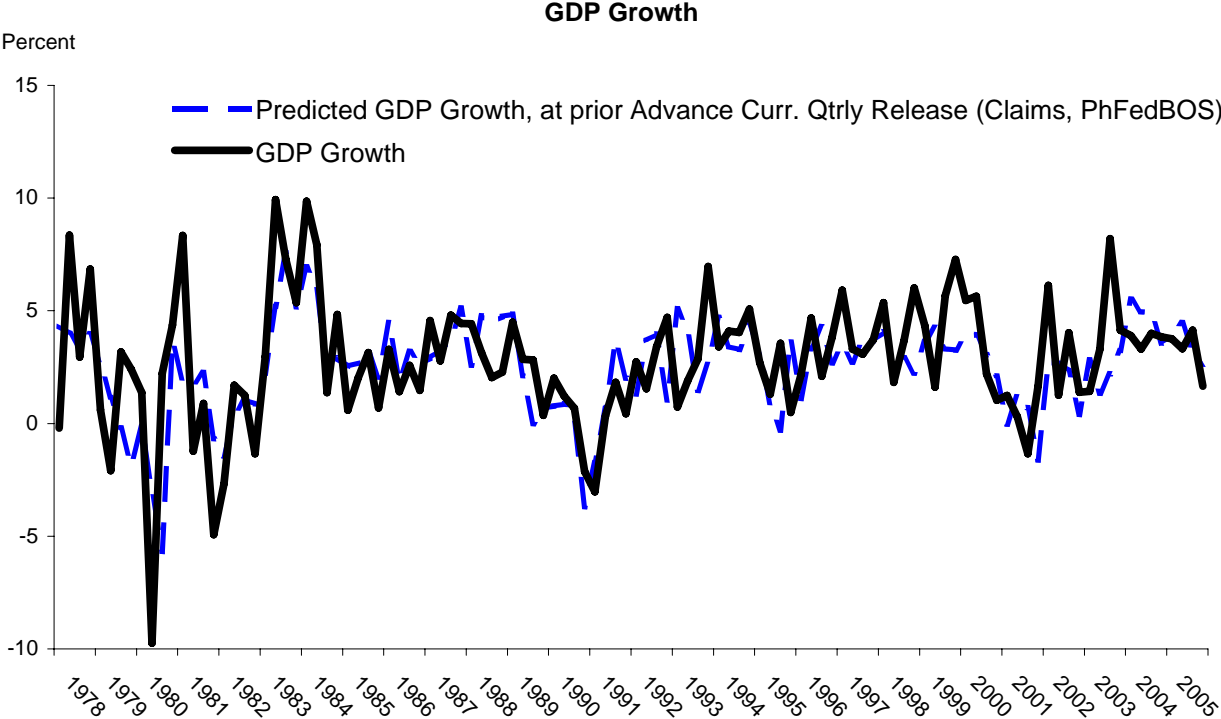


Figure 5: Predicted and Actual Real "Final" Current Quarterly Data





Figure 6: Predicted and Actual Real "Final" Current Quarterly Data





Figure 7: Smoothed Probabilities of Recession

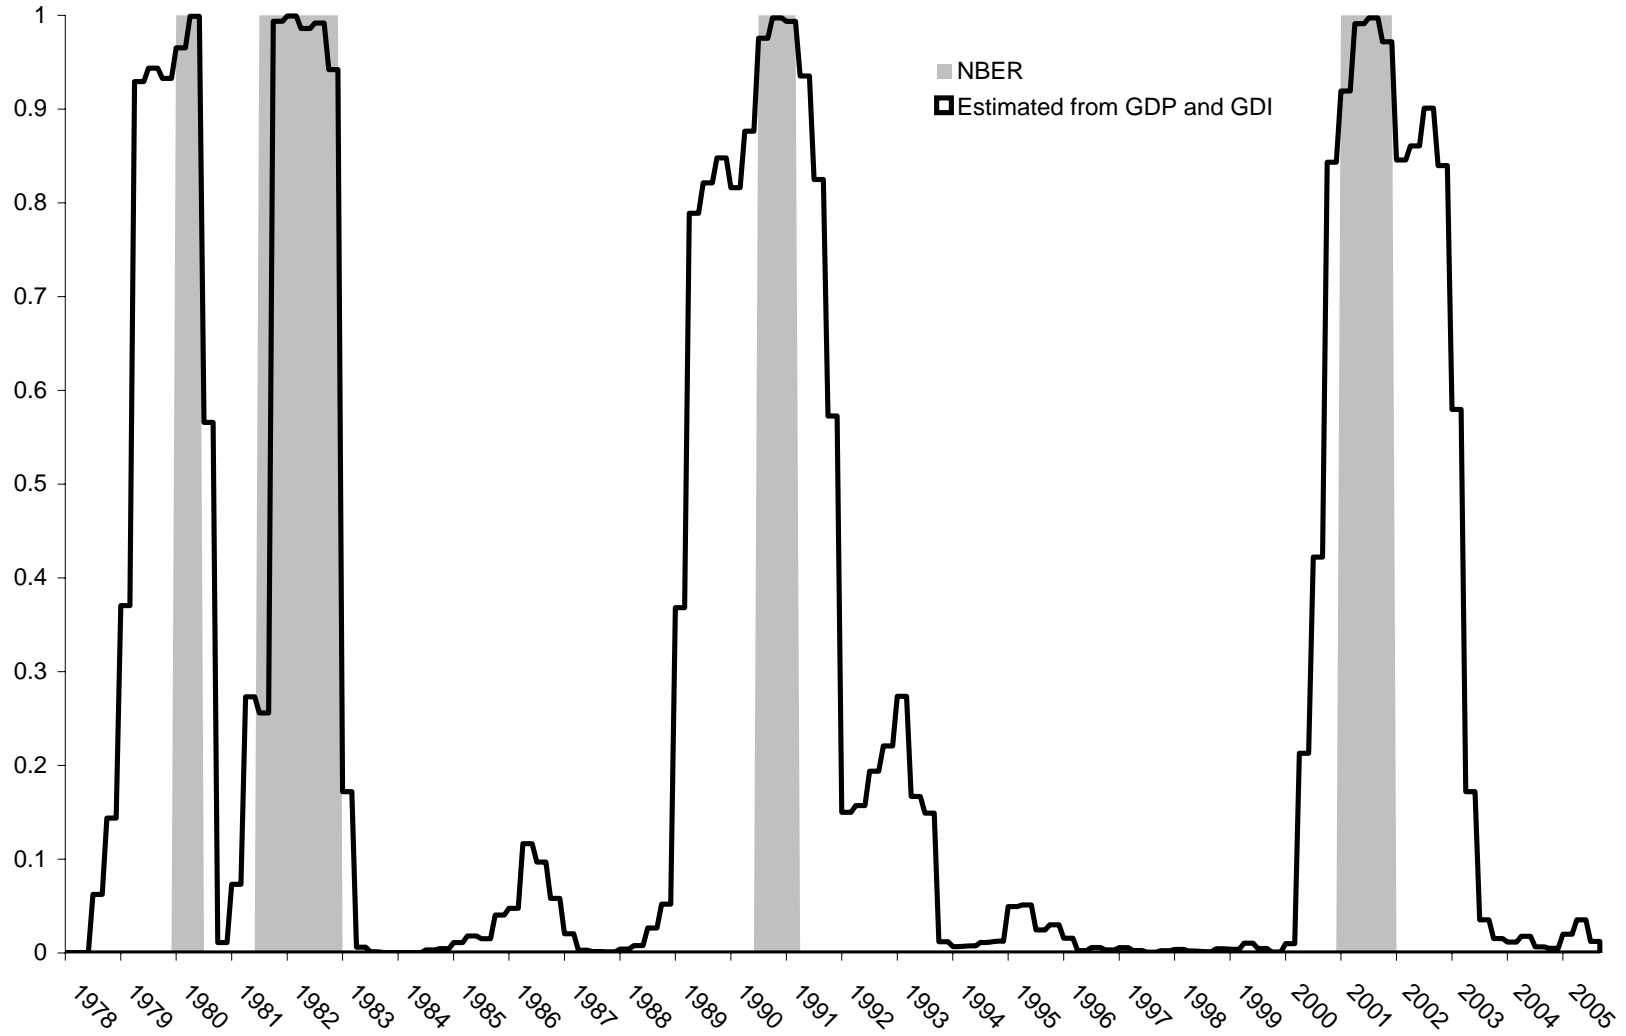
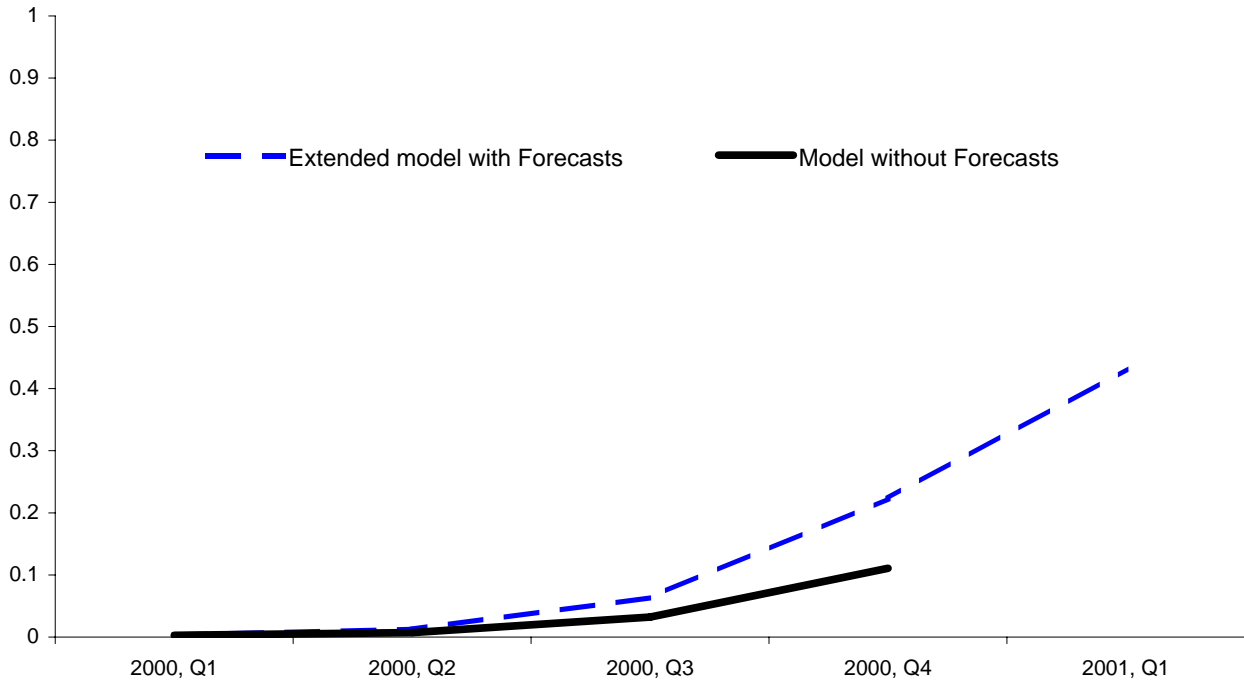


Figure 8: Smoothed Probabilities of Recession

"Advance" 2000Q4 BEA data release, January 2001



"Final" 2000Q4 BEA data release, March 2001

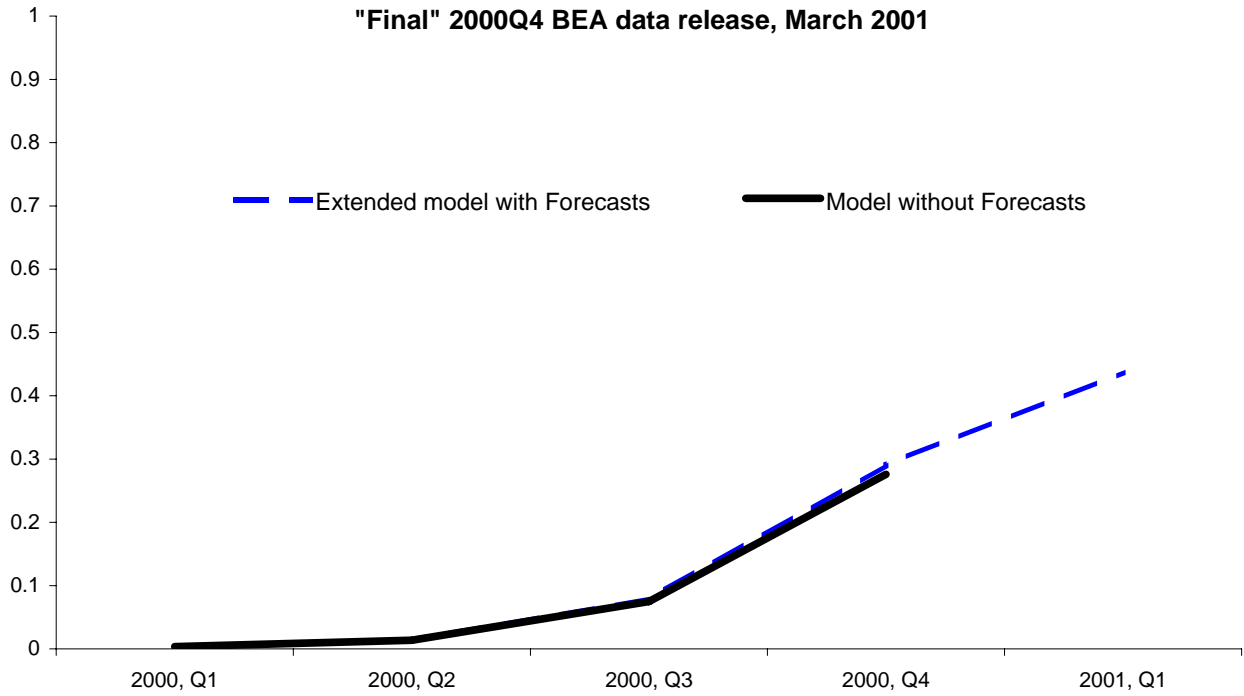
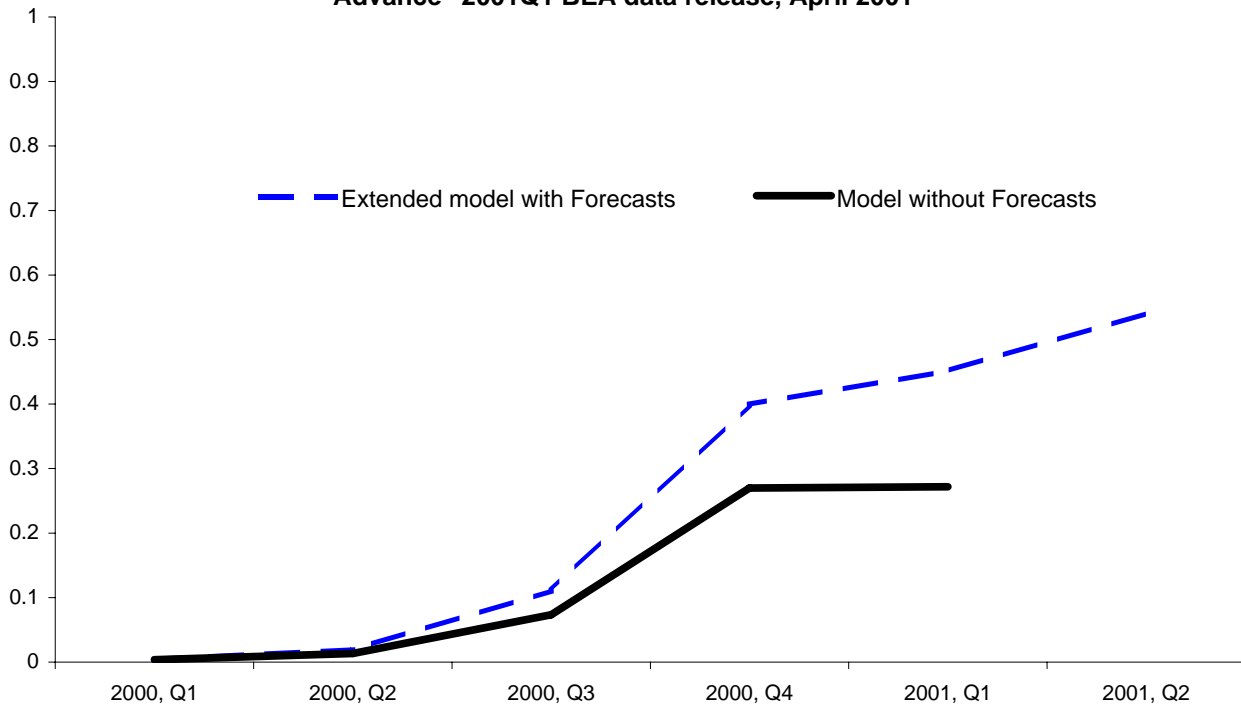


Figure 9: Smoothed Probabilities of Recession
"Advance" 2001Q1 BEA data release, April 2001



"Final" 2001Q1 BEA data release, June 2001

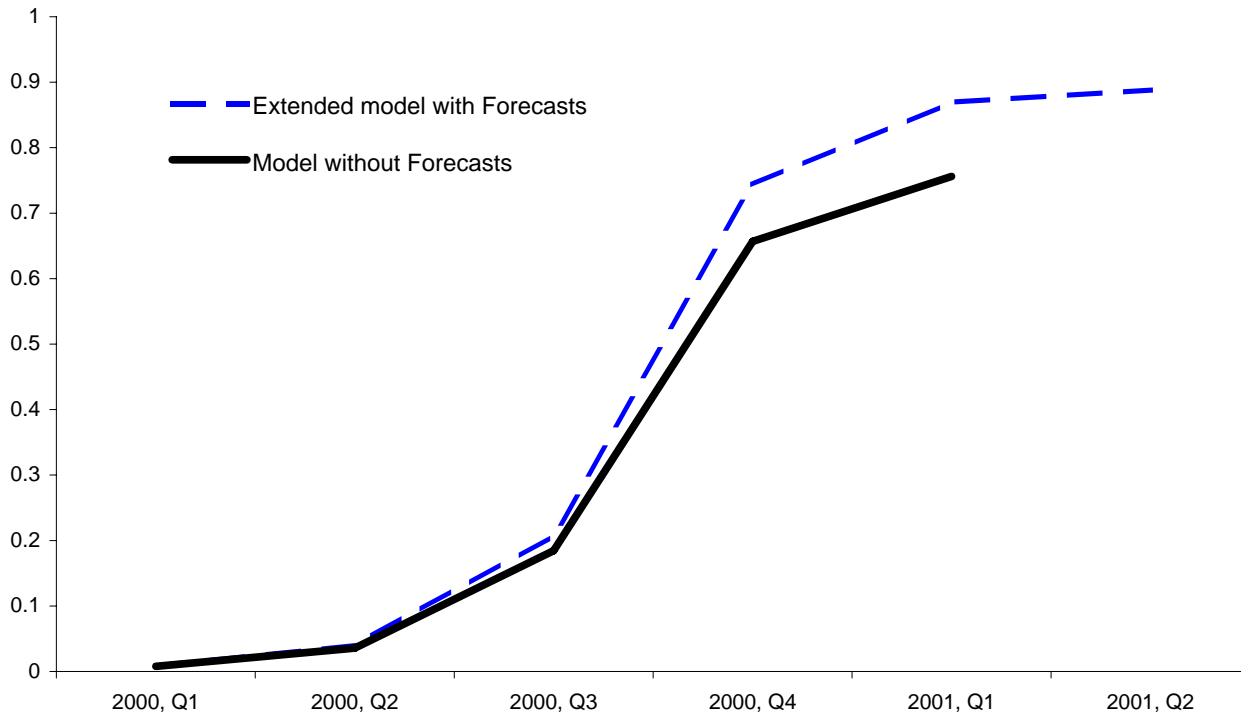


Figure 10: Smoothed Probabilities of Recession

"Final" 2001Q2 BEA data release, September 2001

