

Calling Recessions in Real Time*

James D. Hamilton

Department of Economics

University of California, San Diego

jhamilton@ucsd.edu

May 3, 2010 revised: May 13, 2010

Abstract

This paper surveys efforts to automate the dating of business cycle turning points. Doing this on a real time, out-of-sample basis is a bigger challenge than many academics might presume due to factors such as data revisions and changes in economic relationships over time. The paper stresses the value of both simulated real-time analysis— looking at what the inference of a proposed model would have been using data as they were actually released at the time— and actual real-time analysis, in which a researcher stakes his or her reputation on publicly using the model to generate out-of-sample, real-time predictions. The immediate publication capabilities of the internet make the latter a realistic option for researchers today, and many are taking advantage of it. The paper reviews a number of approaches to dating business cycle turning points and emphasizes the fundamental trade-off between parsimony— trying to keep the model as simple and robust as possible— and making full use of available information. Different approaches have different advantages, and the paper concludes that there may be gains from combining the best features of several different approaches.

JEL classification: E32

*I am grateful for assistance and suggestions provided by Maximo Camacho, Marcelle Chauvet, Jeremy Nalewaik, Gabriel Perez-Quiros, Jeremy Piger, James Stock, Mark Watson, and Jonathan Wright.

1 Introduction

U.S. real GDP has grown significantly over time, today standing at a level more than seven times that seen in 1947 (see Figure 1). Economists broadly agree on the three factors responsible for that long-run growth. First, the population has increased over time, and more people can produce a greater quantity of goods and services. Second, the stock of equipment and facilities that people have to work with has also increased more than six-fold over this period. And third, the techniques of production, such as higher-yielding crops, faster computers, and more efficient management have all produced tremendous improvements in productivity. These three factors—population, capital stock, and technology—are widely regarded to be the main drivers of long-run growth.

But GDP does not increase every single year, and it is of substantial interest to understand why. None of these three factors offer very appealing explanations for downturns. In a recession, we do not lose population; rather, an ever-growing number of people report they are unable to find jobs. Capital stock is not destroyed, but instead sits idle. And although the hypothesis that recessions may be caused by an exogenous decline in productivity has been popular with real business cycle theorists, many of us do not find that account compelling. Something other than these three factors seems to be governing aggregate economic behavior during particular identifiable episodes. These episodes are referred to as economic recessions, indicated by the shaded regions in Figure 1.

The importance of recessions is even more dramatic when we consider a series such as the U.S. unemployment rate in Figure 2. There is no clear trend in this series over the last half century. But there is a quite dramatic tendency for unemployment to rise sharply during periods characterized as economic recession. The habit of accompanying graphs like Figures 1 and 2 with recessions indicated as shaded regions is quite ingrained in the economics profession, and with good reason. One of the first things we want to know about any series is how its fluctuations are

related to movements in and out of economic recession.

The question discussed in this paper is, where should those shaded regions be drawn, that is, when did each particular recession begin and end? The conventional answer for U.S. data is to use the dates determined by the Business Cycle Dating Committee of the National Bureau of Economic Research. This is a group of distinguished scholars who meet periodically to discuss recent data, and from time to time issue a proclamation that a recession began or ended at a particular date in the past. Typically, once the Committee makes a pronouncement, that date is not subsequently revised.

That is a fine approach to use, and this paper does not advocate abandoning it. I nevertheless investigate whether it is possible to supplement the deliberations of the Committee with mechanical algorithms that could process incoming data in real time and issue a judgment on a purely objective basis. There are three potential benefits to doing so. The first is timeliness. The bottom panel of Figure 3 notes the particular dates at which the Business Cycle Dating Committee issued its most recent announcements. These announcements often came long after the fact. For example, the NBER dated the 1990-91 recession as beginning in August 1990 and ending in March 1991. It made the announcement that the recession had begun in April 1991— one month after the NBER later decided that the recession was already over. The end of the 2001 recession was announced in July 2003, which is 28 months after the recession is now deemed to have ended.

A second benefit of a purely objective algorithm for making these determinations in real time is that it would ensure that the process is completely apolitical. Although no one has accused the NBER of altering its announcements on the basis of political considerations, the pressure is undeniably present to delay the announcement that a recession has begun or accelerate the announcement that a recovery has begun if one's goal were to help the incumbent. In the October 13, 1992 debates, then vice-presidential-candidate Al Gore referred to the "worst economic performance since the Great Depression." Although the recession of 1990-91 is now seen as one of the shortest and mildest of the postwar recessions, it was not until December 22, 1992— after

the presidential elections— that the NBER announced that the recession had actually ended in March of 1991. If it had been the case that the recession that began in August 1990 was still ongoing as of October 1992, Gore’s statement would have factual support in the sense that the ongoing recession would have been the longest on record since 1933. Being able to issue these announcements on the basis of a pre-determined, purely mechanical algorithm would insulate the procedure from any possible charges of partisanship.

Third, coming up with a mechanical means to recognize business cycle turning points offers further elucidation of what we actually mean when we say that the economy is in a recession. If the dates assigned by the NBER represent the answer, what is the question? The whole process seems to presuppose that there are some very different factors operating on the economy at some times relative to others, and that these changes have observable implications. Mechanization of the dating procedure can help clarify exactly how and why we assign the dates that we do.

In this paper I survey some of the approaches to dating business cycle turning points, with a focus on algorithms that various analysts have publicly relied on with real-time data. I begin in Section 2 with a review of some early efforts and why the task might be harder than it looks. Section 3 describes in detail a procedure that I have been relying on for the last five years. A number of alternative approaches are discussed in Section 4.

2 What’s so hard?

One’s first thought might be that it shouldn’t be too hard to do better than the multi-year lags sometimes associated with NBER announcements, and indeed the really exciting thing would be to predict business cycle turning points before they occur. What’s so hard about that?

The first answer is that if people could have predicted the recession, it probably wouldn’t have happened. Firms would not have been stuck with inventories, labor, and capital they turned out not to need, and the Federal Reserve probably would have chosen to ease its policy stance

sooner. Economists are used to viewing magnitudes such as stock prices as difficult or impossible to predict if the market is functioning properly. It may be that economic recessions by their nature imply similar fundamental limitations for forecasting.

Second, the data available at the time are often sending a different signal from what one sees once the data are subsequently revised. For example, the top panel of Figure 4 displays real GDP growth rates for each quarter of 2001 as they were reported at the end of 2002. This vintage of data shows 3 successive quarters of declining real GDP, which certainly sounds unambiguously like a recession. On the other hand, the bottom panel shows data for the same quarters as they were actually reported on January 30, 2002. At that date the recession was already over, but recognizing those GDP numbers available at the time as signalling a recession is obviously a much bigger challenge.

A third factor making it difficult to recognize business cycle turning points in real time is the fact that key economic relations are continually changing over time. I illustrate these difficulties by reviewing the real-time track record of two prominent efforts to predict business cycle turning points.

2.1 Stock and Watson's original business cycle index model.

One of the impressive early efforts along these lines came from the business cycle indicators developed by James Stock and Mark Watson (1989, 1991). Their coincident index model postulated that an observed vector of monthly indicators \mathbf{y}_t was related to an unobserved scalar c_t thought to represent the business cycle according to

$$\mathbf{y}_t = \mathbf{k} + \gamma(L)c_t + \mathbf{u}_t$$

$$\mathbf{D}(L)\mathbf{u}_t = \boldsymbol{\varepsilon}_t$$

$$\phi(L)c_t = \delta + \eta_t.$$

The observed variables in \mathbf{y}_t consisted of growth rates of industrial production, personal income, sales, and employment. White noise disturbances ε_{rt} and η_t were taken to be mutually uncorrelated and $\mathbf{D}(L)$ is a diagonal matrix in the lag operator L . Thus the model postulated that the observed dynamics of each y_{rt} could be explained in terms of idiosyncratic autoregressive terms u_{rt} and common dependence on the business cycle c_t . This will be recognized as a state-space model with state vector $(\mathbf{u}'_t, c_t)'$ and observation vector \mathbf{y}_t , for which algorithms for estimating the unknown parameters by maximum likelihood and forming an optimal inference about the unobserved state of the business cycle c_t are well known. Stock and Watson further defined an unobserved variable S_t to be unity if the economy was in recession at date t . They interpreted a recession to be a particular pattern followed by c_t ,

$$S_t = 1 \text{ if } \{c_{t-j}\}_{j=0}^8 \in B_t,$$

with the set B_t designed so as to best mimic historically-assigned NBER recession dates. The basic Kalman filter algorithms allow one to form an inference about current or past recessions

$$P(S_{t-h} = 1 | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots, \mathbf{y}_1; \hat{\boldsymbol{\theta}}) \tag{1}$$

for $h = 0, 1, 2, \dots$, or forecast of future recessions

$$P(S_{t+h} = 1 | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots, \mathbf{y}_1; \hat{\boldsymbol{\theta}}). \tag{2}$$

For the latter purpose Stock and Watson also developed a leading-indicator generalization of the model of the form

$$\mathbf{y}_t = \mathbf{k} + \boldsymbol{\gamma}(L)c_{t-1} + \boldsymbol{\Gamma}(L)\mathbf{y}_{t-1} + \boldsymbol{\varepsilon}_t$$

$$c_t = \delta + \alpha(L)c_{t-1} + \beta(L)'y_{t-1} + \eta_t$$

for y_t now a vector based on levels or growth rates of some conventional leading indicators along with a few new indicators proposed by Stock and Watson.

The in-sample performance of their model is summarized in Figure 5. The top panel gives the contemporaneous inferences ($h = 0$ in equation (2)), which looked quite good. The model might seem to have produced a “false positive” in 1967, though this was a widely recognized economic slowdown which perhaps could reasonably have been categorized as an economic recession. The 3- and 6-month ahead forecasts (lower panels in Figure 5) also looked extremely promising.

Encouraged by these in-sample results, Stock and Watson began in 1988 to report updated probabilities of recession each month using the latest data. Initially these were distributed by FAX and mail to assorted researchers and members of the press, and later were posted on the web.¹

Figure 6 displays the model’s real-time track record over the 1990-91 recession as described by Stock and Watson (1993). The leading index proved to be a disappointment— the recession of 1990 came and went, with the model always predicting that no recession was coming (see bottom panels of Figure 6). The contemporaneous index in fact did by November of 1990 signal that a recession had started (top panel), but the model thought it was going to be sufficiently short-lived that the 3- and 6-month-ahead probabilities always remained below 50% even though ex-post probabilities eventually recognized that a recession had arrived at some earlier date.

What went wrong? One of the intriguing new leading indicators that Stock and Watson discovered was the spread between the yield on commercial paper and Treasury bills. Figure 7 shows that this spread had spiked up dramatically prior to each of the recessions in their original sample, but did very little out of the ordinary in the 1990-91 recession for which their model was

¹ An archive of releases for 1999 to 2003 is available at <http://www.economics.harvard.edu/faculty/stock/files/xri.zip>.

on real-time display.²

Stock and Watson subsequently released probabilities from both their original leading index and an alternative leading index that did not make use of interest rates or interest rate spreads. The models' performance over the 2001 recession was similar to that for the 1990-91 downturn. The alternative leading index calculated a 6-month-ahead probability of recession $P(t + 6|t)$ that remained below 20% throughout 2001, while $P(t + 6|t)$ as calculated from the original leading index peaked at 35% for $t =$ October 2001. Nonetheless, inferences from the coincident index again did reasonably well. On July 3, 2001, Stock and Watson reported $P(t|t) = 76\%$ for $t =$ May 2001, and the value reached 97% for $t =$ October 2001. In retrospect, if Stock and Watson had emphasized using their approach to recognize a turning point some time after it had been crossed— that is, having most faith in the historical inferences (1)— its record would have looked pretty good. But the goal of forecasting the moves in advance proved too ambitious.

2.2 Using the yield curve to predict turning points.

Many of the subsequent academic papers in this area have focused on procedures that might be simpler and more robust. Quite a few academic studies have suggested that the slope of the yield curve— for example, the spread between yields on a 10-year Treasury bond and a 3-month Treasury bill shown in the top panel of Figure 8— seems to be extremely promising as a predictor of recessions; see among others Estrella and Mishkin (1998), Chauvet and Potter (2005), Kauppi and Saikkonen (2008), and Katayama (2009). As the yield curve again became inverted in August of 2006, many observers began to wonder whether this meant another recession would soon arrive.

One aspect of the situation that made this a difficult call at the time was the fact that although the 3-month rate was above the 10-year rate, the overall level of the 3-month rate was lower than it had been prior to any recession since 1960 (see the bottom panel of Figure 8). Many observers argued that the low overall level of interest rates mitigated somewhat the recessionary signal given

² Notwithstanding, credit spread indicators remain a very promising variable for forecasting real economic activity, as demonstrated by the recent analysis by King, Perlie and Levin (2007) and Gilchrist, Yankov, and Zakrajšek (2009).

by the inverted yield curve.

One specification that came to be used in ongoing real-time announcements was one developed by Jonathan Wright (2006). Wright defined a historical indicator H_t to be unity if it subsequently proved to be the case that NBER would declare any of the following four quarters to have been characterized by an economic recession:

$$H_t = \begin{cases} 1 & \text{if } S_{t+1} = 1, S_{t+2} = 1, S_{t+3} = 1, \text{ or } S_{t+4} = 1 \\ 0 & \text{otherwise} \end{cases} .$$

His model then calculated the probability of a recession occurring some time over the next year from

$$P(H_t = 1 | i_{10y,t}, i_{3m,t}, i_{ft}) = \Phi(-2.17 - 0.76(i_{10y,t} - i_{3m,t}) + 0.35i_{ft})$$

for $\Phi(\cdot)$ the $N(0, 1)$ cumulative distribution function, $i_{10y,t}$ the 10-year Treasury yield, $i_{3m,t}$ the 3-month Treasury yield, and i_{ft} the fed funds rate.

The website Political Calculations reported weekly updates of the predictions of this model between April 2006 and August 2008 (see Figure 9). The highest this probability ever reached was 50% on April 4, 2007, after which the probability monotonically declined. It stood at only 10% on August 20, 2008, right before one of the worst 6 months that the economy has experienced over the last half century.

These two examples illustrate that a good in-sample fit is no guarantee of out-of-sample real-time performance. As Yogi Berra observed, it's tough to make predictions, especially about the future. In this paper I instead pursue the more modest goal of trying to recognize a turning point soon after it occurred using algorithms that would be robust with respect to data revisions and out-of-sample structural change. The next section describes one very simple approach that so far seems to have a reasonable track record.

3 Real-time inference based on GDP alone

Consider the following formulation of the fundamental question: what is different about the behavior of GDP during quarters that the NBER characterized as recession compared with those characterized as expansion? Chauvet and Hamilton (2006) answered this by collecting U.S. GDP growth rates from the 45 quarters between 1947:Q2 and 2004:Q2 that the NBER ended up describing as part of an economic recession. This subsample of 45 observations has a mean growth rate of -1.2% (expressed at an annual rate) and standard deviation of 3.5. The top panel of Figure 10 plots a nonparametric estimate of the density of this subsample³. The remaining 184 expansion quarters had a mean of 4.5 and standard deviation of 3.2, with density given by the middle panel. The observed sample of 225 observations can be viewed as a mixture of these two distributions, with 20% coming from the recession distribution and the remaining 80% from the expansion distribution, as shown in the bottom panel of Figure 10.

Suppose we are given only the value for a given quarter's GDP growth, and aren't told how the NBER is eventually going to characterize that quarter. Can we infer from which of the two distributions it is most likely to have been drawn? Let $S_t = 1$ if quarter t is eventually declared by the NBER to be part of a recession, $S_t = 2$ if it is eventually declared to be part of an expansion, and let y_t denote the observed GDP growth. Let $P(S_t = 1, y_t)$ denote the joint probability that quarter t is a recession and has GDP growth given by y_t ,

$$P(S_t = 1, y_t) = f(y_t|S_t = 1)P(S_t = 1),$$

which is found by multiplying the height of the curve in the top panel of Figure 10 by 0.2. Likewise $P(S_t = 2, y_t)$ is obtained by multiplying the height of the middle curve by 0.8. The optimal inference,

$$P(S_t = 1|y_t) = \frac{P(S_t = 1, y_t)}{P(S_t = 1, y_t) + P(S_t = 2, y_t)},$$

³ This was calculated using the DENSITY command in RATS with a Gaussian kernel and bandwidth of 3.

will then be recognized as simply the ratio of the height of the green curve to the height of the black curve in the bottom panel of Figure 10. For example, if we observe GDP growth of $y_t = -6$ we could be quite confident this would be characterized as a recession, whereas if we observe $y_t = +6$, it will almost surely be classified as an expansion.

Such a rule seems to satisfy the requirements of being extremely simple and robust. But unfortunately, it is not terribly helpful, because the vast majority of observations will fall in a range where they don't send a very clear signal. But there is a second feature of the NBER dates that can be quite helpful—the value of S_t is pretty likely to be the same as S_{t-1} . Ninety-five percent of the observations for which $S_{t-1} = 2$ it was the case that S_t also equalled two, whereas 78% of the observations for which $S_{t-1} = 1$ were followed again by $S_t = 1$. Thus, even if y_t alone does not give us much of a useful signal, the value of y_{t-1} can help us refine it. For example, an observed value of $y_t = 1.0$ tells us very little, implying (by calculating the ratio of heights just described) a probability of recession $P(S_t = 1|y_t = 1.0) = 0.21$ that is little different from the unconditional odds. But suppose this follows after a quarter of strong growth, $y_{t-1} = 6.0$, from which the height ratio calculation yields a very low probability of having been in recession the previous quarter: $P(S_{t-1} = 1|y_{t-1} = 6.0) = 0.074$. Before seeing the quarter t GDP numbers, we would accordingly have known that the probability of a recession observation in period t is only 10% rather than the unconditional probability of 20%:

$$P(S_t = 1|y_{t-1} = 6.0) = 0.074(0.78) + (1 - 0.074)(1 - 0.95) = 0.10.$$

Having observed $y_{t-1} = 6.0$, to use the period t value for y_t we'd then want to weight the blue curve in the top panel by 0.1 rather than by 0.2 to arrive at the appropriate conditional probability:

$$P(S_t = 1|y_t, y_{t-1}) = \frac{f(y_t|S_t = 1, y_{t-1})P(S_t = 1|y_{t-1})}{f(y_t|S_t = 1, y_{t-1})P(S_t = 1|y_{t-1}) + f(y_t|S_t = 2, y_{t-1})P(S_t = 2|y_{t-1})}.$$

From this calculation it turns out $P(S_t = 1|y_t = 1.0, y_{t-1} = 6.0) = 0.11$; having seen strong growth the preceding quarter, we'd be reasonably confident that the expansion is continuing this quarter despite the weak GDP report. We can of course extend such calculations, doing better

yet using the whole history of observations up to the present to calculate $P(S_t = 1|y_t, y_{t-1}, \dots, y_1)$. For that matter, the same principles allow us to use observations on $y_{t+1}, y_{t+2}, \dots, y_T$ to refine our assessment of where the economy was at some date t , which we refer to as a full-sample smoothed probability:

$$P(S_t = 1|y_T, y_{T-1}, \dots, y_1).$$

I described above how such an inference could be implemented nonparametrically with little assumed structure other than some very broad features of the NBER dating record. But there's little reason we need to be nonparametric, since the distribution in the top panel of Figure 10 could clearly be well approximated with a Gaussian density. Also, little is lost by assuming that the variance of the distributions in the top two panels is the same. In other words, we could parameterize the top two panels of Figure 10 as

$$y_t|S_t = 1 \sim N(\mu_1, \sigma^2)$$

$$y_t|S_t = 2 \sim N(\mu_2, \sigma^2),$$

expecting to use $\mu_1 = -1.2$, $\mu_2 = 4.5$, and $\sigma = 3.4$. If we let π denote the unconditional probability that a given quarter will be characterized by recession (expecting $\pi = 0.2$), we can characterize the height of the black curve in the bottom panel of Figure 10 for observation $t = 1$ as

$$f(y_1; \mu_1, \mu_2, \sigma, \pi) = \pi\phi(y_1; \mu_1, \sigma) + (1 - \pi)\phi(y_1; \mu_2, \sigma)$$

for $\phi(\cdot)$ the Normal density, with the proposed inference given by

$$P(S_1 = 1|y_1) = \frac{\pi\phi(y_1; \mu_1, \sigma)}{\pi\phi(y_1; \mu_1, \sigma) + (1 - \pi)\phi(y_1; \mu_2, \sigma)}.$$

For observation $t = 2$ we were proposing to calculate

$$P(S_2 = 1|y_2, y_1) = \frac{\xi_2\phi(y_2; \mu_1, \sigma)}{\xi_2\phi(y_2; \mu_1, \sigma) + (1 - \xi_2)\phi(y_2; \mu_2, \sigma)}$$

for

$$\xi_2 = p_{11}P(S_1 = 1|y_1) + (1 - p_{22})P(S_1 = 2|y_1)$$

with p_{11} denoting the probability of a recession continuing, $p_{11} = P(S_t = 1|S_{t-1} = 1)$, and p_{22} the probability of an expansion continuing. Iterating for $t = 3, 4, \dots$, we have

$$P(S_t = 1|y_t, y_{t-1}, \dots, y_1) = \frac{\xi_t \phi(y_t; \mu_1, \sigma)}{\xi_t \phi(y_t; \mu_1, \sigma) + (1 - \xi_t) \phi(y_t; \mu_2, \sigma)} \quad (3)$$

for

$$\xi_t = p_{11}P(S_{t-1} = 1|y_{t-1}, y_{t-2}, \dots, y_1) + (1 - p_{22})P(S_{t-1} = 2|y_{t-1}, y_{t-2}, \dots, y_1).$$

If we interpreted the persistence of recessions and expansions as the sole source of serial dependence in observed GDP growth rates, we could also use the above recursion to calculate a conditional likelihood function for the t th observation:

$$\begin{aligned} f(y_t|y_{t-1}, y_{t-2}, \dots, y_1) &= \sum_{j=1}^2 f(y_t|S_t = j)P(S_t = j|y_{t-1}, y_{t-2}, \dots, y_1) \\ &= \xi_t \phi(y_t; \mu_1, \sigma) + (1 - \xi_t) \phi(y_t; \mu_2, \sigma). \end{aligned}$$

Using the unconditional probability given by $\pi = (1 - p_{22})/(1 - p_{22} + 1 - p_{11})$, we could then calculate the sample log likelihood as

$$\sum_{t=1}^T \log f(y_t|y_{t-1}, y_{t-2}, \dots, y_1; \boldsymbol{\theta}) \quad (4)$$

solely as a function of observed GDP growth rates and the assumed parameter vector $\boldsymbol{\theta} = (\mu_1, \mu_2, \sigma, p_{00}, p_{11})'$, and maximize this log likelihood with respect to $\boldsymbol{\theta}$.

Chauvet and Hamilton (2006) used data on GDP growth rates over 1947:Q2 to 2004:Q2 to maximize (4) with respect to $\boldsymbol{\theta}$. The maximum likelihood estimates are reported in the last column of Table 1. Note that these estimates $\hat{\boldsymbol{\theta}}$ are based solely on observed GDP and make no use of the historical NBER classifications. It is therefore of considerable interest that the MLE $\hat{\boldsymbol{\theta}}$ turns out to be quite close to what we would have expected based on the discussion above of NBER classifications. The exercise suggests that the separation of the historical record into periods of expansion and contraction is not just an artifact created by the NBER Business Cycle Dating Committee, but instead is exactly the same kind of classification one would see in the GDP

data themselves if one took the perspective that they come from an mixture of two distributions with different means and transitions between the two regimes governed by an unobserved Markov chain.⁴

The top panel of Figure 11 plots the filter inference (3) for each date t in the 1947:Q2-2004:Q2 sample, along with NBER-determined recession dates. The correspondence with the NBER dates is quite close, suggesting that this algorithm offers a promising alternative for recognizing business cycle turning points. However, although the formula (3) only makes use of data through date t , it would be a mistake to suppose that it would do as well in a real-time setting, for two reasons. First, the value used for the parameter vector $\hat{\theta}$ was in fact based on the full sample of observations through 2004:Q2. Second and more importantly, the value one would be using for any historical observation y_t in that sample was the number reported in the 2004:Q2 vintage set, which could differ substantially from the number as initially released. To try to gauge the usefulness as a real-time tool, for each date t , Chauvet and Hamilton (2006) used a data set for GDP as it would have been reported at that date (obtained from the Croushore and Stark (2003) data base maintained by the Federal Reserve Bank of Philadelphia) both to estimate the parameter vector θ as well as to form an inference. Use of data as actually available at the time resulted in a considerable deterioration of reliability. We therefore recommended waiting one quarter for data to be revised and to be able to use the added precision afforded by one-quarter smoothing before making a declaration, so that one only tries to form an inference about S_t when the GDP growth rate for the next quarter, y_{t+1} , is first released. This real-time one-quarter-smoothed inference is plotted in the lower panel of Figure 11. This exercise suggested that such an inference might indeed offer a useful method for recognizing business cycle turning points.

In our 2006 paper we recommended the following decision rule. When the one-quarter

⁴ One could of course propose a more elaborate parametric model, such as allowing further dynamics in GDP growth beyond those induced by the business cycle, as in the original specification of Hamilton (1989). However, even as fundamental an indicator as GDP has had changing business-cycle behavior over time (e.g., McConnell and Perez-Quiros, 2000). The extreme simplicity of the model described here may give it more robustness with respect to these changes than more detailed specifications.

smoothed inference $P(S_t = 1|y_{t+1}, y_t, \dots, y_1; \hat{\theta}_{t+1})$ first exceeds 0.65, declare that a recession has started, and at that time assign a probable starting point for the recession as the beginning of the most recent set of observations for which $P(S_{t-j} = 1|y_{t+1}, y_t, \dots, y_1; \hat{\theta}_{t+1})$ exceeds 1/2. The recession call would remain in effect until $P(S_t = 1|y_{t+1}, y_t, \dots, y_1; \hat{\theta}_{t+1})$ falls below 0.35, and at that time assign a probable ending point for the recession as the beginning of the most recent set of observations for which $P(S_{t-j} = 1|y_{t+1}, y_t, \dots, y_1; \hat{\theta}_{t+1})$ is less than 1/2.

The starred entries in Table 2 summarize the simulated real-time performance of this algorithm. The dates assigned to historical recessions are quite close to those determined by the NBER, and would typically be announced at about the same time as the NBER announcements. On the basis of these results, I was persuaded that this approach offered a promising alternative for categorizing business-cycle turning points on a real-time basis. Since July of 2005, I have been regularly reporting at the website www.econbrowser.com the value of $P(S_t = 1|y_{t+1}, y_t, \dots, y_1; \hat{\theta}_{t+1})$ each quarter when the advance GDP numbers are first released. Figure 12 plots these values as a function of t . To the left of the vertical line are what I call “simulated real-time calculations”—for each t , the value plotted is the number one would have calculated with the release of quarter $t+1$ GDP numbers as they were initially released at the time. To the right of the vertical line are what I call “actual real time”, meaning that the number plotted for each quarter t is exactly the value that was publicly reported on the day that the $t+1$ numbers were first released. The index has continued to perform well in terms of its true out-of-sample track record. The algorithm calculated that the 11th postwar U.S. recession began in the fourth quarter of 2007, agreeing with the NBER’s determination that 2007:Q4 marked the business cycle peak. The algorithm’s announcement came in January 2009, one month later than the NBER’s declaration in December 2008. The algorithm determined in April 2010 that the recession ended in 2009:Q2. As of the time of this writing, the NBER has not yet declared an end to the most recent recession.

Although the algorithm has so far proved to be robust in terms of its out-of-sample performance, it was not particularly timely in making the call that the recession had started, coming

one month later than the NBER’s declaration, and much later than all observers understood that a powerful economic downturn was underway. Why was the signal from GDP indecisive until January of 2009? Figure 13 shows that the GDP numbers available at the time showed slightly negative growth for 2007:Q4 and 2008:Q3, +0.9% growth for 2008:Q1, and +2.8% for 2008:Q2. Is that convincing evidence that a recession had begun? The probability generated in October 2008 (following release of the 2008:Q3 advance estimate) was 46.9%, consistent with a suspicion that a recession could well have begun, but not definitive. It was not until the 3.9% GDP drop in 2008:Q4 that the index climbed to 88.4%, at which time the benefit of the smoothed inference raised the likelihood above 50% for observations going back to 2007:Q4. Hence the algorithm was unable to declare until January 2009 that a recession had likely started in 2007:Q4.

Given what we were attempting to do with this procedure—form an inference based solely on GDP growth—it is hard to dispute that this was the correct call given the data at the time, and I feel that the procedure performed well in terms of that objective. But it is clearly worth taking a look at the real-time performance of various other procedures that were making use of alternative economic indicators.

4 Real-time track record of models based on other indicators

I motivated the calculations above as an essentially model-free robust inference, while noting that the calculations could alternatively be derived as the optimal statistical inference about an unobserved recession variable S_t that takes on the value of 1 or 2 according to the outcome of a Markov chain and with GDP growth y_t conditional on S_t coming from one of two Normal distributions. As described in Hamilton (1994, Chapter 22), one can also conduct similar inference on much more elaborate parametric formulations by postulating that the history of being in recessions affects the conditional density of some observed vector of indicators \mathbf{y}_t according to $f(\mathbf{y}_t | \mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \dots, \mathbf{y}_1, S_t, S_{t-1}, \dots, S_1; \boldsymbol{\theta})$ for $f(\cdot)$ an arbitrary density governed by unknown pa-

rameters θ . This section reviews the real-time experience with inferring business cycle turning points based on different indicators and alternative Markov-switching time series models.

4.1 Gross domestic income.

The national income accounting convention is that every dollar of value added by definition generates a dollar in income for somebody. We could measure GDP either by looking at expenditures or by looking at income, and should in principle get exactly the same number either way. The Bureau of Economic Analysis does indeed make calculations based on both methods, but in practice the number obtained is not exactly the same, and the difference between gross domestic product and gross domestic income is officially reported by the BEA as a “statistical discrepancy.” If the BEA measurements were all correct, that discrepancy would be zero.

Jeremy Nalewaik (2007, 2010) has argued that GDI may provide helpful information in addition to GDP, particularly for purposes of recognizing business-cycle turning points. The top panel of Figure 14 plots the growth rates of GDP and GDI; note that GDI was sending a significantly more pessimistic signal than GDP during 2007 and 2008. Nalewaik (2007) developed a generalization of the model discussed in the previous section in which GDI and GDP are treated jointly as a bivariate vector indicator. He has been distributing the probabilities calculated by that model the third month of each quarter since 2006:Q3 as part of internal Federal Reserve Board communications. The bottom panel of Figure 14 plots his simulated and actual real-time experience with this method. Nalewaik provided me with these values in September of 2008, at which point his algorithm had generated a strong recession signal for 2007:Q4 through 2008:Q2.

Does this mean that we should base an inference on GDI rather than GDP? By definition, we do not have a good understanding of what the statistical discrepancy may represent, and the conventional view is that GDP may be the slightly more reliable indicator. Certainly for the most recent recession, GDI would have provided a more timely signal, but for some earlier episodes, the GDI-based inference from the bottom panel of Figure 14 is less reliable than the GDP-based

inference in Figure 12. If we have to rely on just one for future unknown incoming data, I am more comfortable sticking with GDP, even though GDI did better last time. It might be possible to try to combine the information from the two series in a more parsimonious way, for example, using a weighted average of the two measures.

4.2 Unemployment rate.

The dramatic cyclical behavior of unemployment seen in Figure 2 suggests we might have better success with a measure based on the unemployment rate rather than GDP, with sharp and rapid increases in unemployment being perhaps a defining characteristic of an economic recession. My 2005 paper proposed the following model for y_t the monthly unemployment rate,

$$y_t = c_{s_t} + \phi_1 y_{t-1} + \phi_2 y_{t-2} + u_t \quad (5)$$

where u_t is distributed Student t with estimated degrees of freedom and S_t is assumed to follow an unobserved 3-state Markov chain. This specification will be recognized as a generalization of the parametric interpretation of the approach to GDP above, in allowing for additional dynamics beyond shifts in the business-cycle phases (as represented by nonzero values for ϕ_1 and ϕ_2) and allowing for three rather than two business cycle states (as represented by the three possible values for the intercept, c_1 , c_2 , and c_3). The three business cycle phases might be interpreted as normal growth ($S_t = 1$), mild recession ($S_t = 2$), and severe recession ($S_t = 3$).

Figure 15 shows the smoothed probability that the U.S. economy was in either the mild or severe recession phase calculated from this model as of early September of 2008. At that point the model was declaring with 95% confidence that a recession had started. The parameter estimates $\hat{\theta}$ were taken from Hamilton (2005) and I publicized this particular graph on September 5, 2008. Does that qualify as out-of-sample, actual real time? Although I (and most other observers) were persuaded that the U.S. had entered a recession at this time, I would be uncomfortable using this model (even though it's my own!) as a basis for issuing a fully automated announcement for purposes of the next cycle. The reason is that although (5) seems to capture nicely the timing

of historical business cycles, it is harder to motivate as a primitive, first-order description of what we mean by a recession in the same way as the GDP-based inference discussed in the previous section. The unemployment rate exhibits changes in cyclical behavior over time, as demographic and other variables evolve significantly. Hence, even though it's worked historically, I'd have less confidence that such a rule will work as well on the next out-of-sample recession. Even so, this kind of inference, like the GDI inference just discussed, offers a very useful supplement to the inference from GDP.

4.3 Multiple monthly indicators.

Chauvet (1998) and Kim and Nelson (1998) proposed incorporating business-cycle shifts directly into a model like that in Stock and Watson (1989). The unobserved state of the business cycle is represented by a scalar c_t that is subject to regime shifts,

$$c_t = \alpha_{S_t} + \phi c_{t-1} + \eta_t \tag{6}$$

where $\eta_t \sim N(0, \sigma_\eta^2)$ and $\alpha_{S_t} = \alpha_1$ when the economy overall is in a recession ($S_t = 1$) and $\alpha_{S_t} = \alpha_2$ in expansion, with transitions between $S_t = 1$ and 2 again governed by a Markov chain. Each of 4 indicators is presumed related to the business cycle according to

$$y_{rt} = \lambda_r c_t + v_{rt} \quad \text{for } r = 1, 2, 3, 4 \tag{7}$$

with the idiosyncratic factor v_{rt} itself exhibiting $AR(1)$ dynamics:

$$v_{rt} = \theta_r v_{r,t-1} + \varepsilon_{rt}. \tag{8}$$

Variants of the same four indicators used by Stock and Watson— growth rates of industrial production, personal income, sales, and employment— can also be used here. Chauvet and Hamilton (2006) and Chauvet and Piger (2008) conducted simulated real-time exercises on these type of models suggesting that they could perform well in real time.

Since August 2006, Piger has been posting the probabilities calculated by this model almost every month on a publicly available web page.⁵ The actual real-time filter inferences for each month, $P(S_t = 1 | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots, \mathbf{y}_1; \hat{\boldsymbol{\theta}}_t)$ are plotted in the top panel of Figure 16. This filter probability first moved above 50% in Piger’s November 1, 2008 report which was based on data describing August. The probability moved to 99.2% when the September data became available but fell back to 16.1% with the next month’s data which showed industrial production and real personal income less transfers to be growing again in October. Chauvet and Piger (2008) had recommended an inference rule of declaring a recession as soon as three consecutive values of $P(S_{t-k} = 1 | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots, \mathbf{y}_1; \hat{\boldsymbol{\theta}}_t)$ were all above 80%. This threshold was passed with the release of the October data, for although $P(S_t = 1 | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots, \mathbf{y}_1; \hat{\boldsymbol{\theta}}_t)$ was only 16.1% based on $t = \text{October}$ 2008 data, $\Pr(S_{t-k} = 1 | \mathbf{y}_t, \mathbf{y}_{t-1}, \dots, \mathbf{y}_1; \hat{\boldsymbol{\theta}}_t)$ was above 80% for $k = 1, 2,$ and 3 (see the bottom panel of Figure 16). Chauvet and Piger’s (2008) announced rule would then date the start of the recession as the earliest k for which this probability was above 50%, which turned out to be February 2008. Thus their approach would have announced on January 1, 2009 that a recession had started in February 2008, although the filter probabilities available in January 2009 raised the possibility that the recession could already have been over by October. Subsequent data confirmed the downturn was ongoing (panels 2 and 3 of Figure 16). The Chauvet-Piger rule of waiting for a reading of 3 consecutive months below 20% would have resulted in a declaration in January 2010 that a recovery likely began in July of 2009.

Marcelle Chauvet has periodically posted⁶ real-time assessments from a closely related specification adapted from Chauvet and Hamilton (2006). One important difference from the specification of Chauvet and Piger is the reliance on the BLS household survey for the measure of employment rather than the establishment payroll data. Figure 17 shows the sequence of smoothed probabilities associated with different vintages of data. This model would have sent a signal in

⁵ Apparently the January 1, 2009 posting was missed due to a holiday data release and personal commitments. I am most grateful to Jeremy Piger for providing me with these data.

⁶ See <http://sites.google.com/site/marcellechauvet/probabilities-of-recession>.

August 2008 that a recession had begun in December 2007. Like Piger’s calculations, it showed a temporary hope of an end when only data through October 2008 were available. Chauvet used this model to announce on her website in October of 2009 that the recession ended in June or July of 2009.⁷

This actual real-time experience confirms that there is useful additional information in the monthly indicators, though some caution is necessary before trying to use this framework as an actual real-time basis for recognizing business cycle turning points. Chauvet and Potter (2010) also emphasize the need to allow for instability over time of monthly indicator models.

4.4 Inference from panel data.

Hamilton and Owyang (2009) proposed a method for forming an inference based on data on employment growth for each of the separate 48 contiguous U.S. states. The assumption is that each state n might be in recession or not for any given quarter t ($S_{nt} = 1$ or 2 for $n = 1, \dots, 48$) with the quarterly employment growth rate for state n and quarter t , y_{nt} , governed by $y_{nt}|S_{nt} \sim N(\mu_{S_{nt}}, \sigma_n^2)$. Our hypothesis was that certain clusters of states might be in recession together and that the determination of which cluster is in recession at any given date was governed by a larger Markov chain. The smoothed probability of each cluster being in recession at any given date is reported in Figure 18. Of particular interest is the blue line, which tracks the probability that we were at that date experiencing a national recession in which every state participated. These track closely the NBER dates, though assign a much longer duration to the recession of 2001 on the basis of its “jobless recovery.”

What is perhaps remarkable is that the model concluded that, with virtual certainty, the U.S. economy had entered a national recession as of 2007:Q2. I publicized this particular graph and its implication that we were well into another recession on Econbrowser on April 3, 2008. Although

⁷ Chauvet has also been calculating probabilities using the payroll employment figures instead of the household survey, and feels that the payroll numbers might better recognize a peak while the household numbers are better for the trough. Based on the payroll numbers, she announced in April 2008 that the recession had started in December.

the inference proved to be prescient, it is again one I wouldn't trust out of sample. For one thing, the model attributed all the comovement between states to these recession clusters, an assumption that was necessary given our estimation algorithm but that is clearly implausible. If you really think you're getting 48 independent observations on something, it's correct that you could be virtually certain of your inference, but this is an artifact of assuming independence rather than a correct assessment of the uncertainty. The fact that employment growth behaved so differently for the 2001 recession as in previous recessions is another reason to be queasy about relying on such a model. It nevertheless again illustrates the potential gains from using multiple indicators, in this case employment data at the level of individual states.

4.5 Mixed-frequency inference.

Camacho and Perez-Quiros (2010) showed how data of different frequencies can be combined in a linear state-space framework. Their model uses observations on Euro-area GDP and employment together with monthly observations on industrial production, retail sales, exports, industrial new orders, Euro-zone Economic Sentiment Indicator, German business climate index, Belgian overall business indicator, and both services and manufacturing purchasing managers indexes for the Euro area. One of the nice contributions of their paper is a very convenient algorithm for processing an unbalanced panel on a real-time daily basis as different data releases or revisions to past data get announced on different days. The forecasts from this model are now being used inside the Eurosystem and so far have an excellent actual real-time track record.

Camacho, Perez-Quiros, and Poncela (2010) adapted the same idea to a latent business-cycle regime model based on a similar set of observed mixed-frequency indicators. Figure 19 reports the probabilities of a Euro-area recession that would have been calculated by their approach based on the actual data available as of each day during 2008 and 2009. The model yields a remarkably sharp and stable inference over this period, with the probability jumping from 3.8% on June 24, 2008 to 98.1% by July 24, 2008. The probability remained above 69% until April 23, 2009, at

which point it fell to 6%. Based on these results, the approach appears extremely promising. There is again of course the concern that with so many estimated parameters, the subsequent out-of-sample real-time performance may deteriorate.

5 Other approaches

I have focused in the last two sections exclusively on Markov-switching time series models, in part because this is the approach with the most clearly established out-of-sample real-time track record. But there are a number of promising ideas from alternative perspectives. Here I very briefly mention a few.

One can think of any probit prediction of NBER classifications, such as that by Wright discussed in Section 2, as being driven by an unobserved latent Gaussian variable z_t^* governed by

$$z_t^* = \alpha + \beta' \mathbf{y}_{t-1} + \varepsilon_t \quad (9)$$

with NBER recession classifications S_t characterized by

$$S_t = 1 \text{ if } z_t^* > 0$$

and therefore optimally forecast using the formula

$$P(S_t = 1 | \mathbf{y}_{t-1}) = \Phi(\alpha + \beta' \mathbf{y}_{t-1}).$$

Dueker (2005) proposed to think of z_t^* as part of a VAR⁸ including also the observed variables \mathbf{y}_t and used Bayesian methods to infer parameters based on historical observations of $\mathbf{y}_1, \dots, \mathbf{y}_T$ and NBER classifications S_1, \dots, S_T , which he called a qual-VAR specification. Given the model

⁸ That is, (9) is generalized to

$$\begin{aligned} z_t^* &= \alpha + \beta'_1 \mathbf{y}_{t-1} + \beta'_2 \mathbf{y}_{t-2} + \dots + \beta'_p \mathbf{y}_{t-p} \\ &\quad + \gamma_1 z_{t-1}^* + \gamma_2 z_{t-2}^* + \dots + \gamma_p z_{t-p}^* + \varepsilon_t \\ \mathbf{y}_t &= \mathbf{c} + \Phi'_1 \mathbf{y}_{t-1} + \Phi'_2 \mathbf{y}_{t-2} + \dots + \Phi'_p \mathbf{y}_{t-p} \\ &\quad + \xi_1 z_{t-1}^* + \xi_2 z_{t-2}^* + \dots + \xi_p z_{t-p}^* + \mathbf{u}_t. \end{aligned}$$

parameters, one can then calculate an optimal forecast of future business cycle phases conditional on whatever is actually known as of some out-of-sample date τ :

$$P(S_{\tau+k} = 1 | \mathbf{y}_\tau, \mathbf{y}_{\tau-1}, \dots, \mathbf{y}_1, S_{\tau-r}, S_{\tau-r-1}, \dots, S_1; \hat{\boldsymbol{\theta}})$$

for $S_{\tau-r}$ the most recent NBER classification known with certainty as of date τ . Dueker (2008) used this model on December 9, 2008 to predict (quite successfully, as it turned out) that the recession would end in July or August of 2009 but that employment growth would not resume until March of 2010.

Leamer (2008) considered instead simple rules of the form

$$S_t = 1 \text{ if } y_{rt} < c_r$$

for y_{rt} some indicator and c_r a threshold chosen to capture historical NBER dates. Leamer proposed using 6-month changes in employment (from both the establishment and household survey), industrial production, and the unemployment rate. In August of 2008 he concluded on the basis of such calculations that:

this algorithm indicates that the data through June 2008 do not yet exceed the recession thresholds, and will do so only if things get much worse.

Subsequent data unfortunately were indeed much worse, and Leamer's indicators produced a clear recession call by October 19 (Hamilton, 2008b).

Berge and Jordà (2009) proposed using the receiver operating characteristic curve as a basis for systematically selecting such indicators y_{rt} , and concluded that the national activity index developed by Stock and Watson (1999) and regularly updated by the Federal Reserve Bank of Chicago is a particularly promising indicator.

Harding and Pagan (2006) formalized the traditional Burns and Mitchell (1946) approach of identifying local peaks and troughs in a series and then looking for maximal correspondence between the inference from individual series to identify an aggregate recession. Chauvet and Piger

(2008) found this approach performed reasonably well with simulated real-time data, though it appeared to be dominated by their 4-indicator Markov-switching specification.

6 Conclusion

This paper discussed approaches to recognizing business cycle turning points using a variety of data sets and specifications. There is a trade-off between parsimony and making full use of all available information. A very simple model based on GDP alone appears reasonably robust to data revisions and structural change and so far has a good real-time out-of-sample track record, but could clearly be improved upon. Averaging the inference from alternative specifications or using Bayesian approaches to constrain more richly parameterized specifications are worth exploring. A particularly promising approach is to combine data of different frequencies.

References

Berge, Travis J., and Òscar Jordà (2009), The classification of economic activity, working paper, University of California, Davis.

Burns, Arthur F., and Wesley C. Mitchell (1946), *Measuring Business Cycles*, New York: NBER.

Camacho, Maximo, and Gabriel Perez-Quiros (2010), Introducing the EURO-STING: Short term indicator of Euro area growth, *Journal of Applied Econometrics*, forthcoming.

Camacho, Maximo, Gabriel Perez-Quiros, and Pilar Poncela (2010), Green shoots in the Euro area: a real time measure, working paper, Bank of Spain, forthcoming.

Chauvet, Marcelle (1998), An economic characterization of business cycle dynamics with factor structure and regime switches, *International Economic Review*, 39:969-996.

Chauvet, Marcelle (2010), The beginning and end of the 2007-2009 recession, <http://sites.google.com/site/crefcus/probabilities-of-recession/The-beginning-and-end-of-the-2007-2009-recession1>, April.

Chauvet, Marcelle, and James D. Hamilton (2006), Dating business cycle turning points, in Costas Milas, Philip Rothman, and Dick van Dijk, eds., *Nonlinear Time Series Analysis of Business Cycles*, pp. 1-54, Amsterdam: Elsevier.

Chauvet, Marcelle, and Jeremy Piger (2008), A comparison of the real-time performance of business cycle dating methods, *Journal of Business Economics and Statistics*, 26: 42-49.

Chauvet, Marcelle, and Simon Potter (2005), Forecasting recessions using the yield curve, *Journal of Forecasting*, 24: 77-103.

Chauvet, Marcelle, and Simon Potter (forthcoming), Monitoring business cycles with structural changes, *International Journal of Forecasting*.

Croushore, Dean, and Tom Stark (2003), A real-time data set for macroeconomists: Does the data vintage matter?, *Review of Economics and Statistics* 85:605-617.

Dueker, Michael (2005), Dynamic forecasts of qualitative variables: A qual VAR model of U.S. recessions, *Journal of Business and Economic Statistics*, 23:96-104.

Dueker, Michael (2008), Predicting the trough and a jobless recovery, http://www.econbrowser.com/archives/2008/12/predicting_the_1.html, Dec. 9, 2008.

Estrella, Arturo, and Frederic S. Mishkin (1998), Predicting U.S. recessions: financial variables as leading indicators, *Review of Economics and Statistics* 80: 45-61.

Gilchrist, Simon, Vladimir Yankov, and Egon Zakrajšek (2009), Credit market shocks and economic fluctuations: evidence from corporate bond and stock markets, *Journal of Monetary Economics*, 56:471-493.

Hamilton, James D. (1989), A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* 57:357-384.

Hamilton, James D. (1990), Analysis of time series subject to changes in regime, *Journal of Econometrics* 45:357-384.

Hamilton, James D. (1994), *Time Series Analysis*, (Princeton University Press, Princeton).

Hamilton, James D. (2005), What's real about the business cycle?, *Federal Reserve Bank of St. Louis Review*, July/August (87, no. 4): 435-452.

Hamilton, James D. (2008a), Regional propagation of business cycles, http://www.econbrowser.com/archives/2008/04/regional_propag.html, April 3.

Hamilton, James D. (2008b), More unhappy numbers, http://www.econbrowser.com/archives/2008/10/more_unhappy_nu.html, October 19.

Hamilton, James D., and Michael T. Owyang (2009), The propagation of regional recessions, working paper, UCSD.

Harding, Donald and Adrian R. Pagan (2006), Synchronization of cycles., *Journal of Econometrics*, 132:59-79.

Katayama, Munechika (2009), Improving recession probability forecasts in the U.S. economy, working paper, Louisiana State University.

Kauppi, Heikki and Pentti Saikkonen (2008), Predicting U.S. recessions with dynamic binary response models, *Review of Economics and Statistics*, 90:777-791.

Kim, Chang-Jin, and Charles R. Nelson (1998), Business cycle turning points, a new coincident index, and tests of duration dependence based on a dynamic factor model with regime-switching, *Review of Economics and Statistics*, 80:188-201.

King, Thomas B., Andrew T. Levin, and Roberto Perli (2007), Financial market perceptions of recession risk, working paper, Federal Reserve Board.

Leamer, Edward E. (2008), What's a recession anyway?, NBER working paper 14221.

McConnell, Margaret M., and Gabriel Perez-Quiros (2000), Output fluctuations in the United States: what has changed since the early 1980's?, *American Economic Review*, 90:1464-1476.

Nalewaik, Jeremy J. (2007), Estimating probabilities of recession in real time using GDP and GDI, working paper, Federal Reserve Board.

Nalewaik, Jeremy J. (2010), The income- and expenditure-side estimates of U.S. output growth, *Brookings Papers on Economic Activity*, forthcoming.

Stock, James H., and Mark W. Watson (1989), New indexes of coincident and leading economic indicators, in Olivier Jean Blanchard and Stanley, Fischer, eds., *NBER Macroeconomics Annual 1989*, 351-394, Cambridge MA, MIT Press.

Stock, James H. and Mark W. Watson (1991), A probability model of the coincident economic indicators, in Kajal Lahiri and Geoffrey H. Moore, eds., *Leading Economic Indicators: New Approaches and Forecasting Records*, (Cambridge University Press, Cambridge, U.K.).

Stock, James H., and Mark W. Watson (1993), A procedure for predicting recessions with leading indicators: econometric issues and recent experience, in James H. Stock and Mark W. Watson, eds., *Business Cycles, Indicators, and Forecasting*, 95-153, Chicago, University of Chicago Press.

Stock, James H., and Mark W. Watson (1999), Forecasting inflation, *Journal of Monetary Economics*, 44:293-335.

Wright, Jonathan H. (2006), The yield curve and predicting recessions, working paper, Federal Reserve Board.

Table 1

Parameter estimates based on (1) characteristics of expansions and recessions as classified by NBER, and (2) values that maximize the observed sample log likelihood of postwar GDP growth rates over 1947:Q2 to 2004:Q2.

Parameter	Interpretation	Value from NBER classifications	Value from GDP alone
μ_1	average growth in expansion	4.5	4.62
μ_2	average growth in recession	-1.2	-0.48
σ	standard deviation of growth	3.4	3.34
p_{11}	prob. expansion continues	0.95	0.92
p_{22}	prob. recession continues	0.78	0.74

Notes to Table 1: Reproduced from Chauvet and Hamilton (2006).

Table 2

Business cycle turning points and dates at which announcements were issued by NBER and the GDP-based algorithm.

Start of recessions				
Peak as determined by NBER	Date NBER made declaration	Recession start as determined by algorithm	Date algorithm made declaration	Algorithm announcement lead (-) or lag (+) in months
1969:Q4	N.A.	1969:Q2	May 1970*	N.A.
1973:Q4	N.A.	1973:Q4	May 1974*	N.A.
1980:Q1	Jun 1980	1979:Q2	Nov 1979*	-7
1981:Q3	Jan 1982	1981:Q2	Feb 1982*	+1
1990:Q3	Apr 1991	1989:Q4	Feb 1991*	-2
2001:Q1	Nov 2001	2001:Q1	Feb 2002*	+3
2007:Q4	Dec 2008	2007:Q4	Jan 2009	+1

Start of expansions				
Trough as determined by NBER	Date NBER made declaration	Recession end as determined by algorithm	Date algorithm made declaration	Algorithm announcement lead (-) or lag (+) in months
1970:Q4	N.A.	1970:Q4	Aug 1971*	N.A.
1975:Q1	N.A.	1975:Q1	Feb 1976*	N.A.
1980:Q3	Jul 1981	1980:Q2	May 1981*	-2
1982:Q4	Jul 1983	1982:Q4	Aug 1983*	+1
1991:Q1	Dec 1992	1991:Q4	Feb 1993*	+2
2001:Q4	Jul 2003	2001:Q3	Aug 2002*	-12
N.A.	N.A.	2009:Q2	Apr 2010	N.A.

Notes to Table 2. Starred entries denote simulated real-time declarations, unstarred are actual real-time delarations. N.A. indicates information is not available.




Figure 1. One hundred times the natural logarithm of U.S. real GDP, 1947:Q1-2010:Q1. Last shaded region covers 2007:Q4-2009:Q2; other shaded regions correspond to NBER recession dates.




Figure 2. Civilian unemployment rate, 1948:M1 to 2010:M3. Seasonally adjusted, from Bureau of Labor Statistics. Last shaded region covers 2007:M12 to 2009:M6; other shaded regions correspond to NBER recession dates.

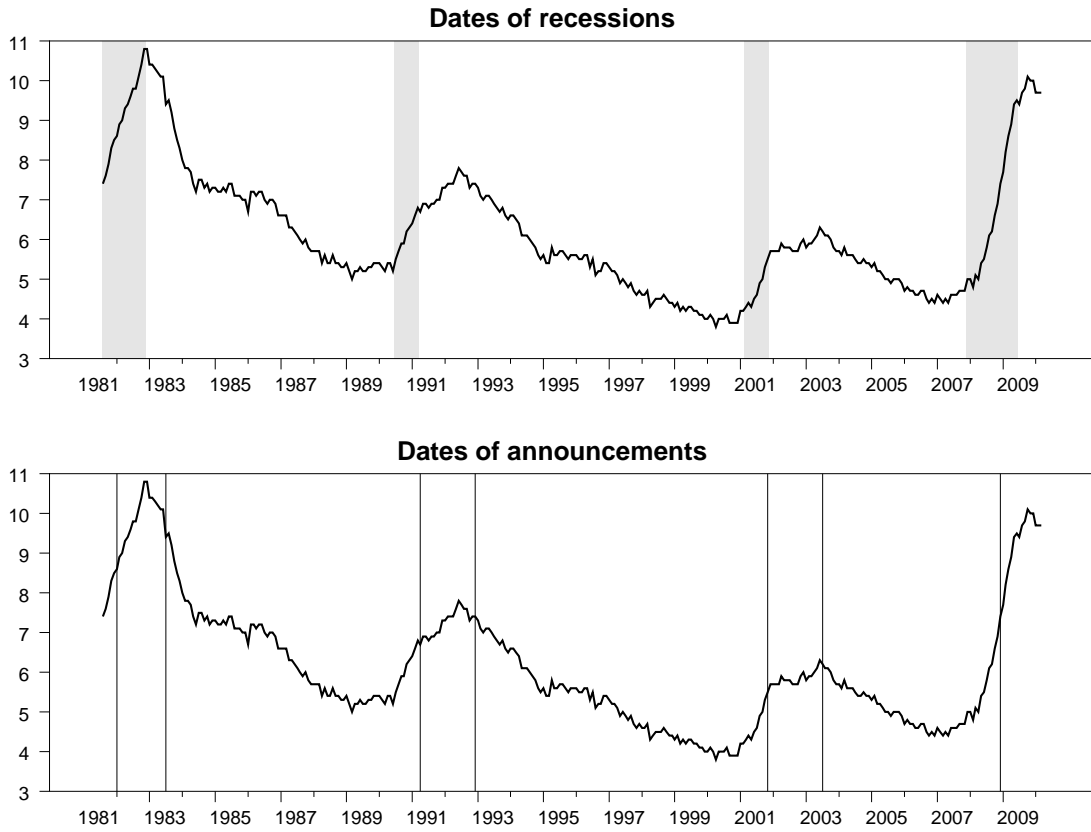


Figure 3. Unemployment, recessions, and dates of NBER announcements. Top panel: same as Figure 2 for 1981:M8 to 2010:M3. Bottom panel: dates at which the NBER announced the business cycle turning points indicated in the top panel.

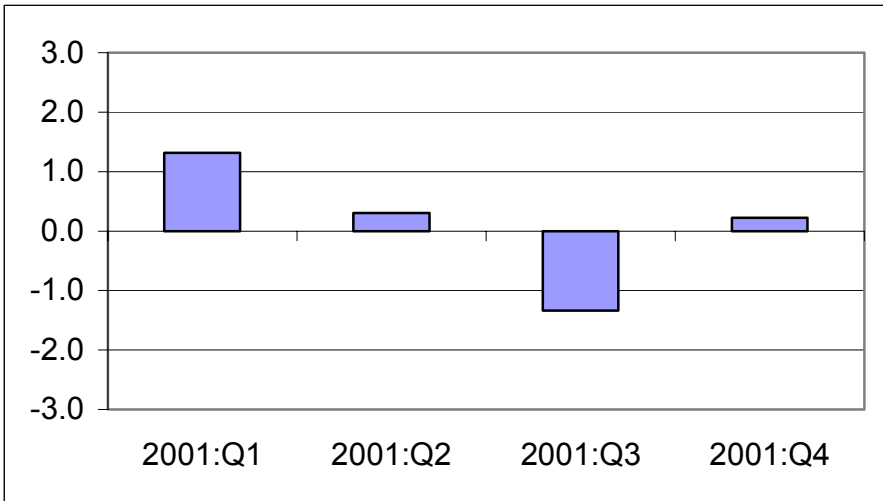



Figure 4. Quarterly growth of real GDP for 2001, quoted at an annual rate. Top panel: as reported on October 31, 2002. Bottom panel: as reported on January 30, Data source: historical data archive of the Federal Reserve Bank of Philadelphia.




Figure 5. In-sample values for probability of recession from Stock-Watson business cycle model. Top panel: contemporaneous probability. Middle panel: 3-month-ahead forecast. Bottom panel: 6-month-ahead forecast. Adapted from Figure 2.1 in Stock and Watson (1993), using code provided by Mark Watson (<http://www.princeton.edu/~mwatson/ddisk/pr.zip>).




Figure 6. Out-of-sample values for probability of recession from Stock-Watson business cycle model. Vertical line drawn for July 1990, which NBER subsequently declared to have been the business cycle peak. Top panel: contemporaneous probability. Middle panel: 3-month-ahead forecast. Bottom panel: 6-month-ahead forecast. Adapted from Figure 2.1 in Stock and Watson (1993), using code provided by Mark Watson (<http://www.princeton.edu/~mwatson/ddisk/pr.zip>).

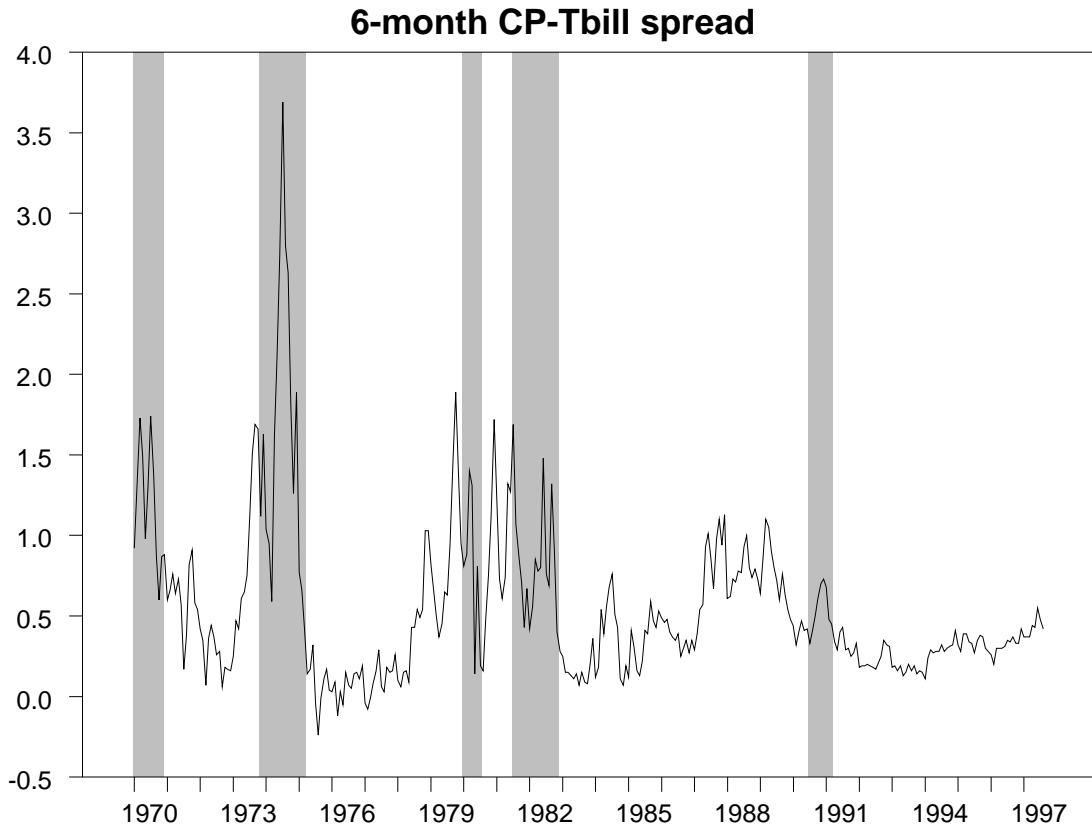


Figure 7. 6-month commercial paper rate minus 6-month Treasury bill rate, 1970:M1 to 1997:M8.

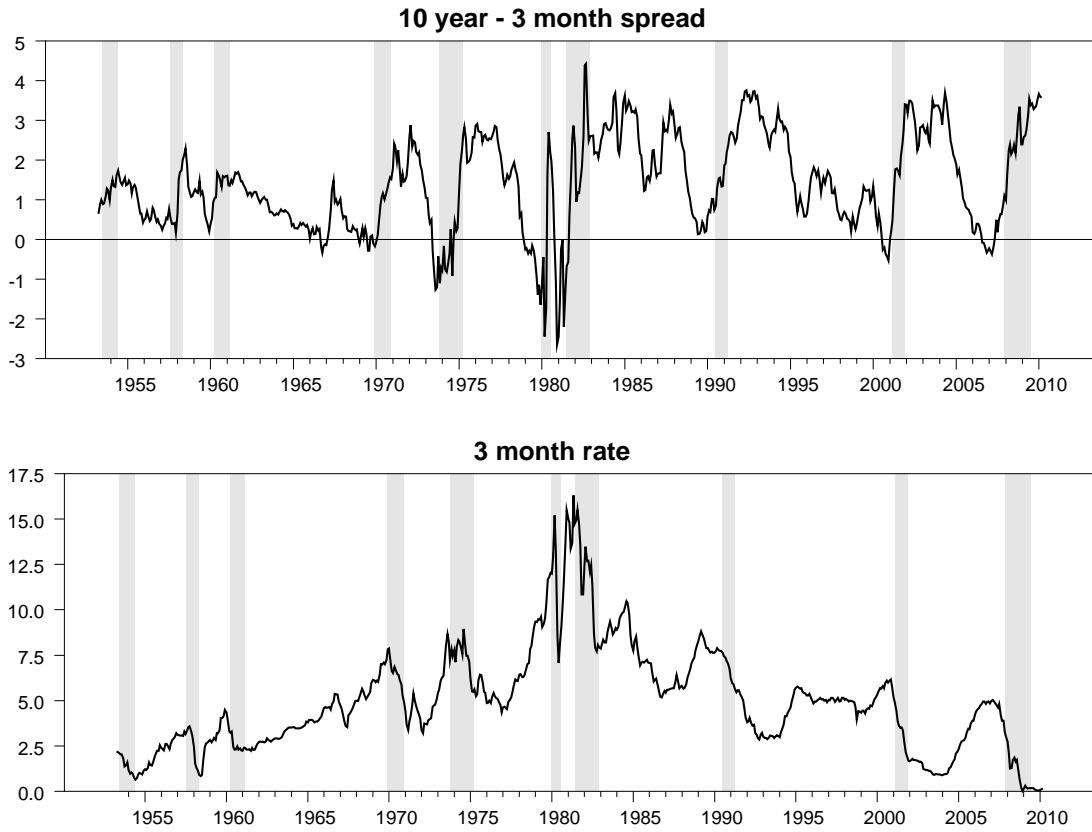
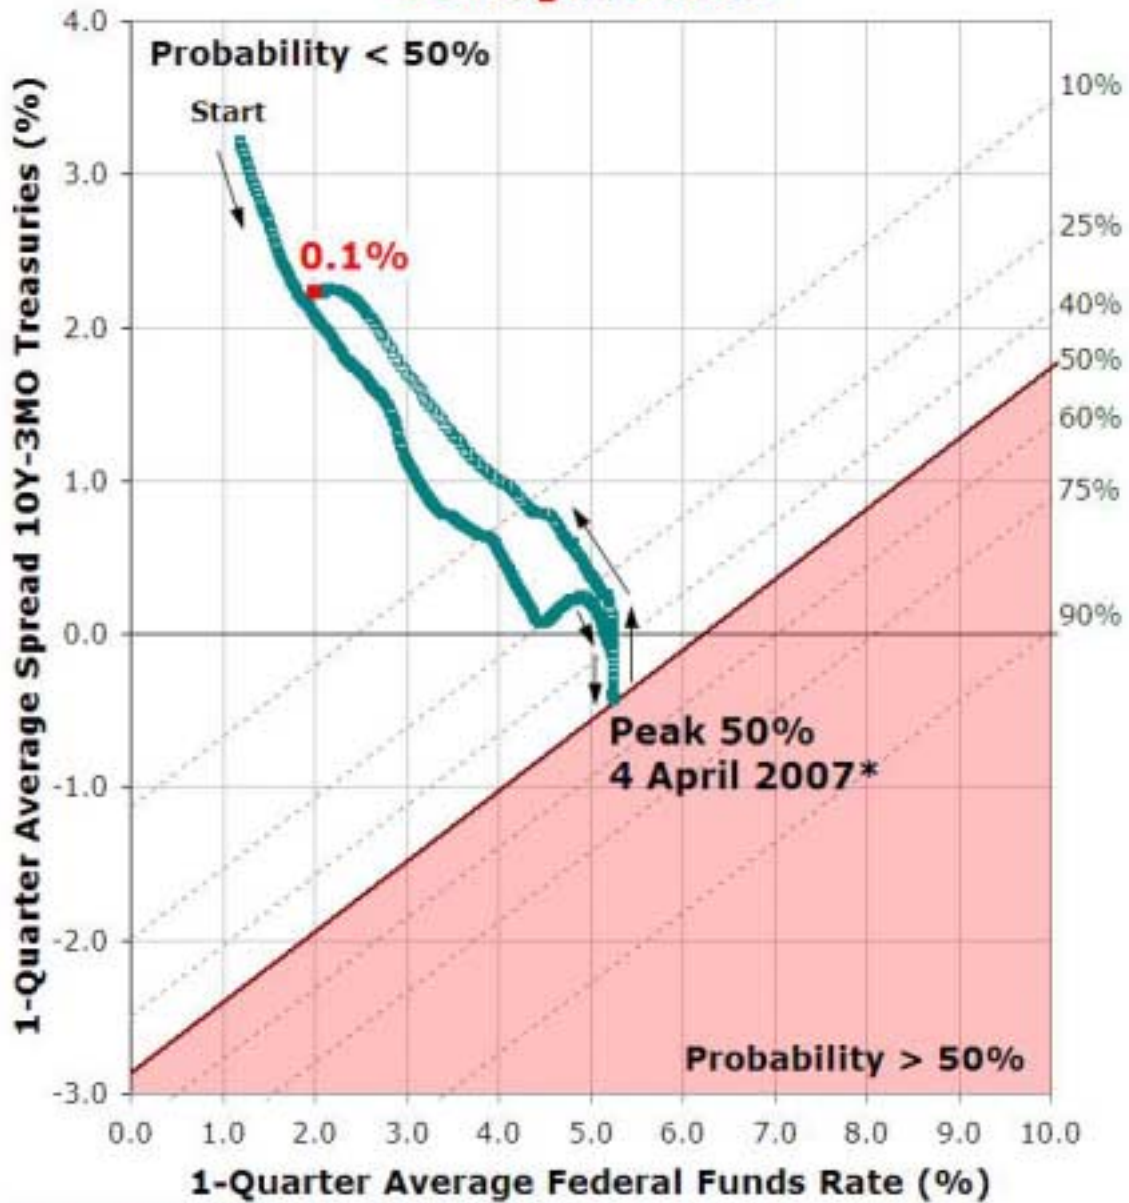


Figure 8. Yield spread and short rate. Top panel: Yield on 10-year Treasury bonds minus that on 3-month T-bills. Bottom panel: level of 3-month T-bill rate.

U.S. Recession Probability Track Starting 20 August 2004 Ending 20 August 2008*



* Recession Probability Shown Applies to Date One Year in Future from Indicated Date

© Political Calculations 2008

Figure 9. Real-time implications of Wright's probit model. Horizontal axis: average fed funds rate over the preceding 90 days. Vertical axis: average 10 year minus 3 month yield over preceding 90 days. Iso-probability implications of the model shown as upward-sloping dashed lines, with circles denoting observed values for each week between August 2004 and August 2008. Source: politicalcalculations.blogspot.com.

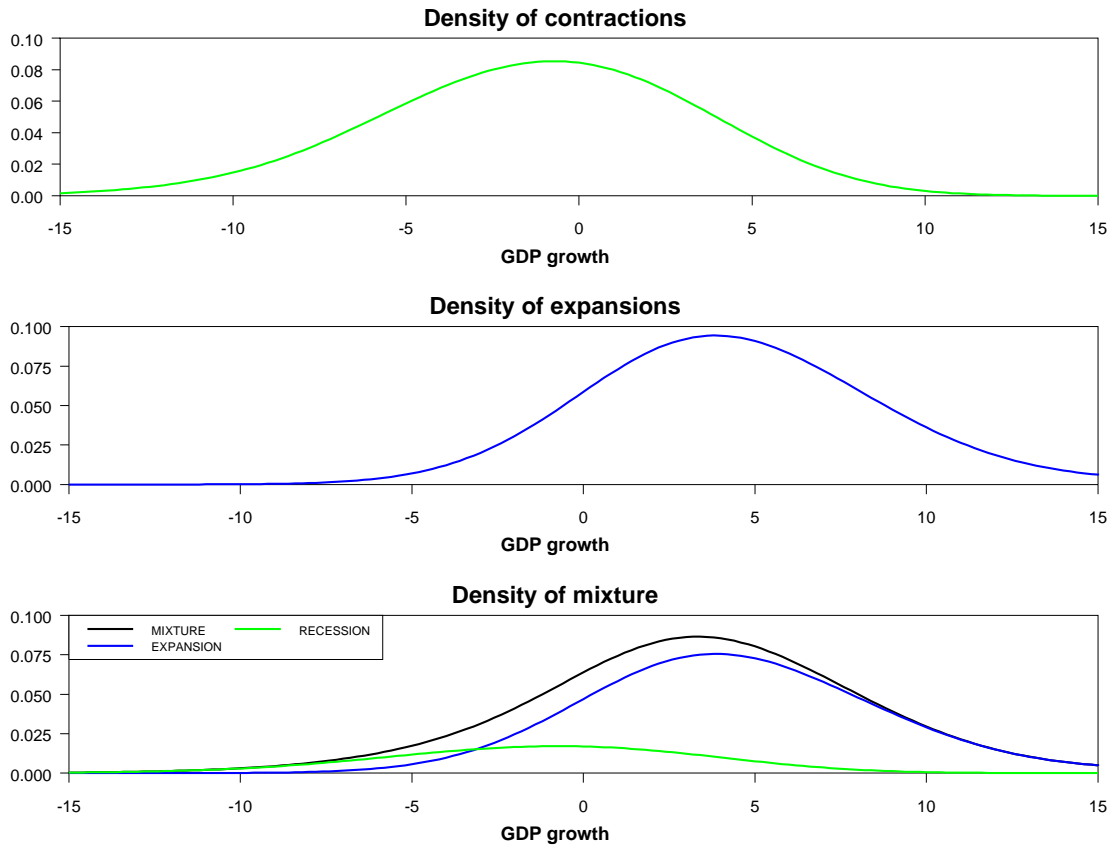


Figure 10. Density of GDP growth in contractions, expansions and overall. Top panel: Density of GDP for quarters characterized by NBER as part of a recession. Middle panel: density for quarters characterized as expansion. Bottom panel: 0.2 times entry in top panel (green) , 0.8 times entry in middle panel (blue) , and sum of the two products (black). Adapted from Chauvet and Hamilton (2006, Figures 2 and 3).




Figure 11. Model-derived recession probabilities. Top panel: filter probability on revised data; plotted series for date t is $P(S_t = 1 | y_{t,T}, y_{t-1,T}, \dots, y_{1,T}; \hat{\theta}_T)$ where $y_{t,T}$ denotes GDP growth for quarter t as reported in vintage released $T = 2004:Q2$, and $\hat{\theta}_T$ is the MLE based on $(y_{1,T}, y_{2,T}, \dots, y_{T,T})$. Bottom panel: one-quarter smoothed probability based on real-time data; plotted series for date t is $P(S_t = 1 | y_{t+1,t+1}, y_{t,t+1}, \dots, y_{1,t+1}; \hat{\theta}_{t+1})$ where $\hat{\theta}_{t+1}$ is the MLE based on $(y_{1,t+1}, y_{2,t+1}, \dots, y_{t+1,t+1})$. NBER-determined recession dates (which were not used in any way to calculate either of the plotted series) are indicated as shaded regions on both graphs. Adapted from Figures 6 and 7 in Chauvet and Hamilton (2006).




Figure 12. GDP-based recession indicator index, 1967:Q4-2009:Q4. Last shaded region covers 2007:Q4-2009:Q2; other shaded regions correspond to NBER recession dates. Prior to 2005, each point on the graph corresponds to a simulated real-time inference that was constructed from a data set as it would have been available four months after the indicated date. After 2005, points on the graph correspond to actual announcements that were publicly released four months after the indicated date.




Figure 13. Quarterly real GDP growth rates for 2007-2008. Quoted at an annual rate, as reported in January 2009.




Figure 14. Real gross domestic product, gross domestic income, and probabilities inferred from joint behavior of GDP and GDP, 1978:Q2 to 2008:Q2. Real GDI calculated as nominal GDP minus statistical discrepancy (BEA Table 1.7.5) divided by implicit GDP deflator. Top panel: black line is quarterly growth rate of real GDP (annual rate), blue line is quarterly growth rate of real GDI. Bottom panel: vertical line at 2006:Q2; values for 2006:Q2 and earlier correspond to simulated real-time probabilities; values for 2006:Q3 and later represent actual real-time probabilities. Source: Jeremy Nalewaik (personal correspondence).

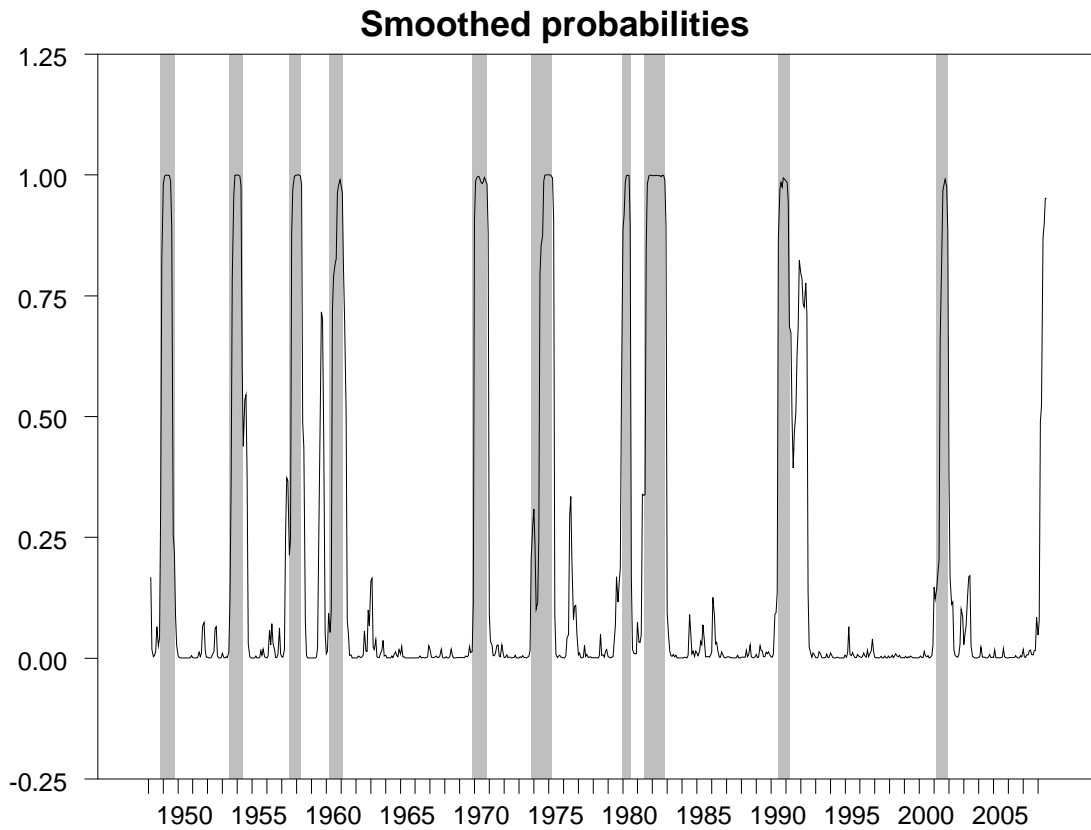


Figure 15. Inference from 3-phase model of unemployment as published on September 5, 2008. Graph shows $P(S_t > 1 | y_T, y_{T-1}, \dots, y_1; \hat{\theta}_T)$ for each t and is reproduced from www.econbrowser.com/archives/2008/09/rising_unemploy.html.

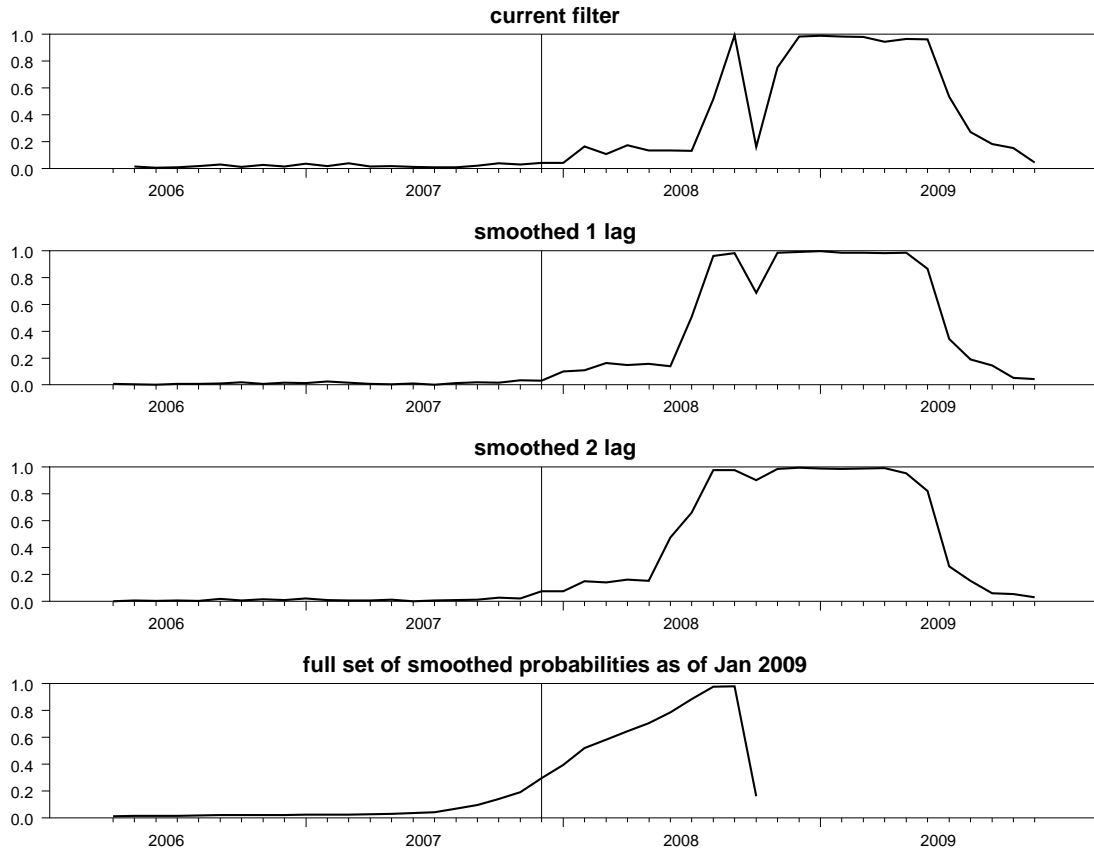


Figure 16. Actual Real-time probabilities from the Chauvet and Piger (2008) 4-indicator model as reported each month at www.uoregon.edu/~jpiger/us_recession_probs.htm. Top panel: filter probability that month t (plotted on horizontal axis) is characterized by recession based on initial release of data describing month t ; value shown would have been actually reported two months after the indicated date. Middle panel: one-month smoothed probability that month t (plotted on horizontal axis) is characterized by recession allowing for 1-month revisions and smoothing; value shown would have been reported three months after the indicated date. Third panel: two-month smoothed probabilities; value shown would have been reported four months after the indicated date. Bottom panel: full sample of smoothed probabilities as of January 2009; values shown would have all been reported on January 1, 2009. Vertical lines drawn at 2007:M12, which NBER declared on December 1, 2008 to have been the business cycle peak.

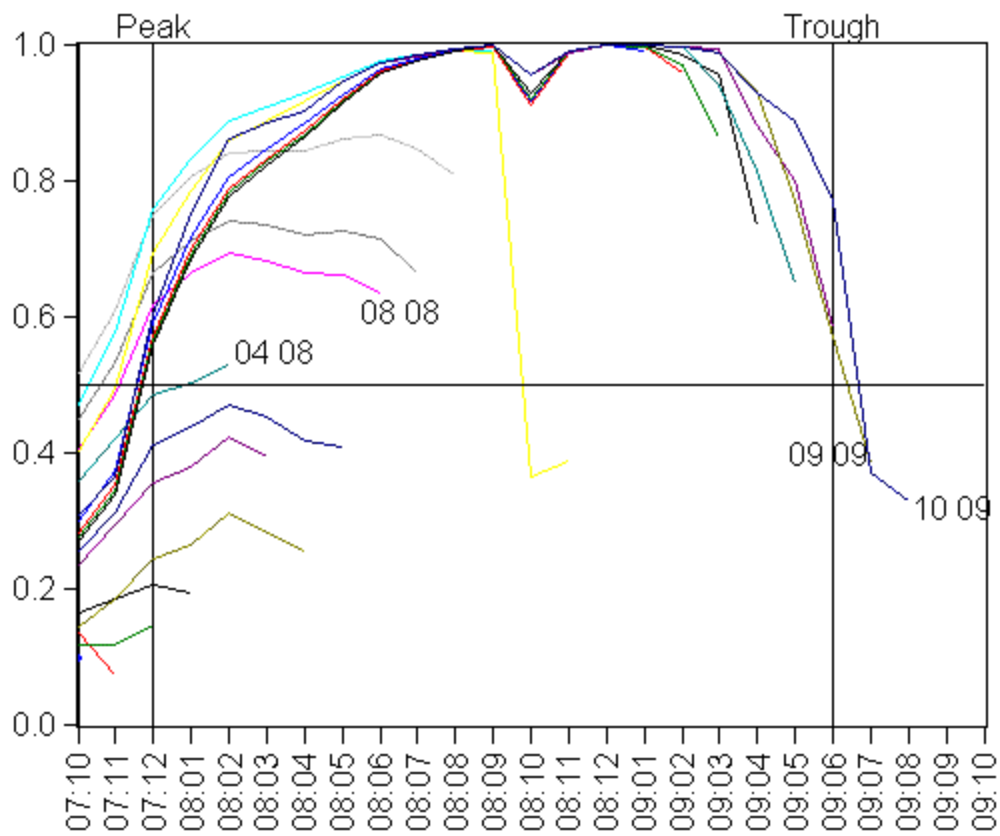


Figure 17. Full sample of smoothed probabilities available for indicated vintage of data as calculated by Chauvet's 4-indicator model. Source: Chauvet (2010).

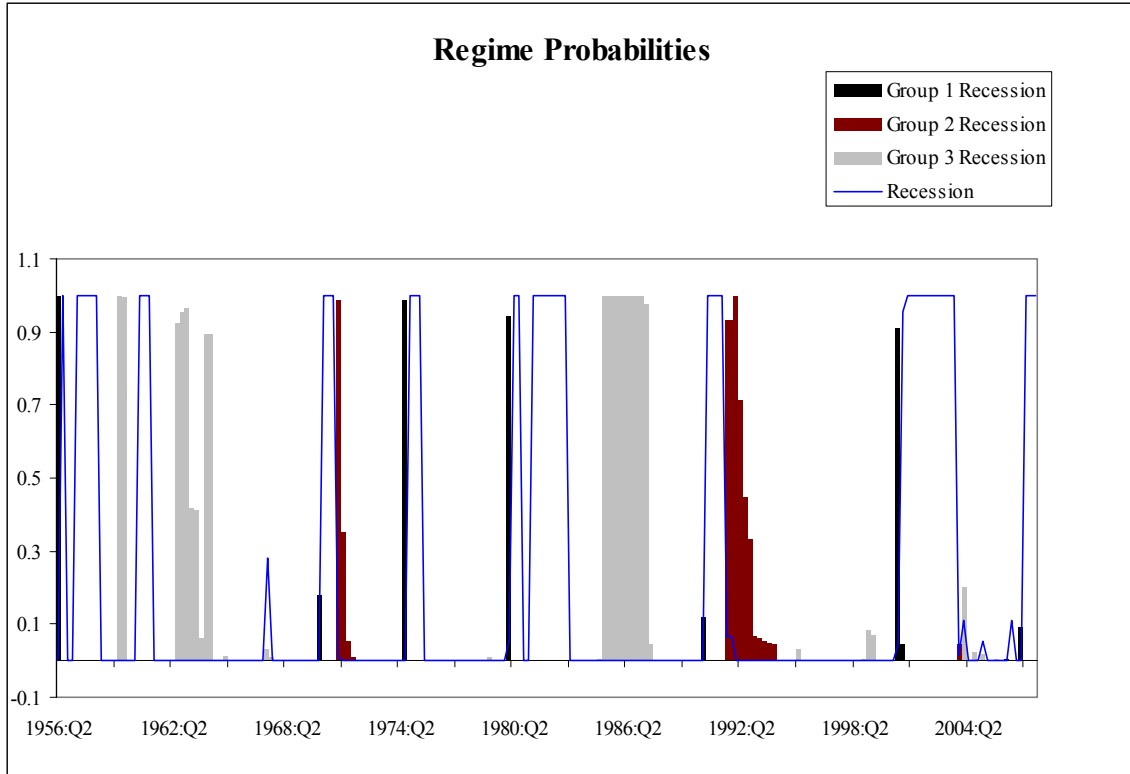


Figure 18. Full-sample smoothed probabilities from the Owyang-Hamilton model of different clusters being in recession, 1956:Q2 to 2007:Q4. Blue lines indicate probability that all states are in recession together, other colors denote conditions when only a subset of U.S. states were in recession. Reproduced from Hamilton (2008a).




Figure 19. Simulated real-time probabilities of Euro-area recession, Jan 1, 2008 to Jan 15, 2010. Value shown for each day is based solely on data actually available as of that day. Source: Camacho, Perez-Quiros, and Poncela (2010).