A Comparison of the Real-Time Performance of Business Cycle Dating Methods

Marcelle CHAUVET
Department of Economics, University of California, Riverside, CA 92521 (*chauvet@ucr.edu*)

Jeremy PIGER
Department of Economics, University of Oregon, Eugene, OR 97403 (*jpiger@uoregon.edu*)

We evaluate the ability of formal rules to establish U.S. business cycle turning point dates in real time. We consider two approaches, a nonparametric algorithm and a parametric Markov-switching dynamic-factor model. Using a new "real-time" dataset of coincident monthly variables, we find that both approaches would have accurately identified the NBER business cycle chronology had they been in use over the past 30 years, with the Markov-switching model most closely matching the NBER dates. Further, both approaches, and particularly the Markov-switching model, yielded significant improvement over the NBER in the speed with which business cycle troughs were identified.

KEY WORDS: Dynamic-factor model; Markov-switching; Recession; Turning point; Vintage data.

1. INTRODUCTION

There is a long tradition in business cycle analysis of separating periods in which there is broad economic growth, called expansions, from periods of broad economic contraction, called recessions. Understanding these phases and the transitions between them has been the focus of much macroeconomic research over the past century. In the United States, the National Bureau of Economic Research (NBER) establishes a chronology of "turning point" dates at which the shifts between expansion and recession phases occur. These dates are nearly universally used in work requiring a definition of U.S. business cycle phases. Since 1978, business cycle dates have been established in real time by the NBER's Business Cycle Dating Committee, which is currently composed of seven academic economists.

The NBER's announcements garner considerable publicity. Given this prominence, it is not surprising that the business cycle dating methodology of the NBER has received some criticism. For example, because the NBER's decisions represent the consensus of individuals who likely bring differing techniques to bear on the question of when turning points occur, the dating methodology is charged as being neither transparent nor reproducible. Also, the NBER has been hesitant to revise business cycle turning point dates, despite the fact that economic data are revised substantially. Finally, the NBER business cycle peak and trough dates are often determined with a substantial lag. For example, the March 1991 and November 2001 business cycle troughs were not announced by the NBER until nearly two years after the fact.

An alternative to the NBER procedures is to use formal rules to date business cycle turning points. Such rules immediately address the first two criticisms above. That is, given that the rules take the form of a formal algorithm or statistical model applied to data, they are both transparent and reproducible. Also, because the rules can be applied to revised data, they provide a straightforward approach to revision of business cycle dates. In this article we evaluate whether such rules can also address the third critique. That is, do these rules provide more timely identification of business cycle dates? Of course, any gain in timeliness must be weighed against any loss of accuracy in establishing the dates. To measure accuracy, we take it as given that the NBER established the correct turning point dates in real time, thus making the NBER chronology the standard for accuracy.

Why are we interested in the speed with which business cycle turning points can be identified? The NBER is likely more concerned with establishing the correct turning point dates than establishing these dates quickly, which breeds additional caution. This caution comes at a low cost if the primary objective is to provide a historical record of business cycle phases. However, as there is substantial evidence that interesting economic dynamics and relationships vary over business cycle phases, economic agents are likely also interested in real-time monitoring of whether a new phase shift has occurred. In this article we provide some formal evidence regarding the speed with which such real-time monitoring can reveal a new turning point in economic activity.

We compare two popular business cycle dating methods, both of which are multivariate in that they use information from many time series to establish business cycle dates. The first is a nonparametric algorithm, developed and discussed in Harding and Pagan (2006) and denoted MHP, for multivariate Harding–Pagan, hereafter. The MHP algorithm proceeds by first identifying turning points as local minima and maxima in the level of individual time series. Next, economy-wide turning points are established by finding dates that minimize a measure of the average distance between that date and the turning points in individual series.

The second approach is a parametric dynamic factor time series model that captures expansion and recession phases as unobserved regime shifts in the mean of the common factor. The unobserved state variable controlling the regime shifts is modeled as following a Markov process as in Hamilton (1989). This Markov-switching dynamic factor model (DFMS), as developed in Chauvet (1998), produces a probability that the economy is in an expansion or a recession at any point in time. These probabilities can then be used to establish turning point dates

Journal of Business & Economic Statistics
January 2008, Vol. 26, No. 1
DOI 10.1198/073500107000000296

using a rule for converting probabilities into a zero/one variable defining which regime the economy is in at any particular time.

We apply these two approaches to a new "real-time" dataset of the four coincident economic variables highlighted by the NBER in establishing turning point dates: (1) nonfarm payroll employment, (2) industrial production, (3) real manufacturing and trade sales, and (4) real personal income excluding transfer payments. In particular, the dating methods are applied as if an analyst had been using them to search for new turning points each month beginning in November 1976, where the data used is the vintage that would have been available in that month. This real-time dataset was collected for this article and has not yet been applied in any other analysis.

The results of this exercise suggest that both approaches are capable of identifying turning points in real time with reasonable accuracy. That is, the first time these methods declare a turning point, the chosen date is usually close to that established by the NBER. The most accurate performance is given by the DFMS model, which provides turning point dates in real time that are usually within one month, and never more than two months, from the corresponding NBER date. Both methods achieve this performance with no instances of "false positives," or turning point dates that were established in real time, but did not correspond to a NBER turning point date. Further, both approaches improve significantly over the NBER in the speed at which business cycle troughs are identified. In particular, the DFMS model would have identified the four business cycle troughs in the sample an average of 249 days, or roughly 8 months, ahead of the NBER announcement, whereas the MHP algorithm would have led by an average of 166 days, or about 5.5 months. However, neither approach provides a corresponding improvement in the speed with which business cycle peaks are identified. Overall, these results suggest that formal dating rules are a potentially useful tool to be used for real-time monitoring of business cycle phase shifts.

Our article makes several contributions to an existing literature on this topic. Layton (1996) evaluated the performance of Markov-switching models of the U.S. coincident index for establishing business cycle turning points. Layton used a "pseudo" real-time analysis in which fully revised data are used in recursive estimations to evaluate the real-time performance of the business cycle dating algorithm. The new real-time dataset we use here provides a more realistic assessment of how the dating rules would have performed, as it does not assume knowledge of data revisions that were not available at the time the rule would have been used. Chauvet and Piger (2003) used real-time data to evaluate the business cycle dating performance of univariate Markov-switching models of employment and real GDP, and Chauvet and Hamilton (2006) did a similar exercise for multivariate Markov-switching models. These articles consider only Markov-switching models, whereas here we compare Markov-switching models to nonparametric algorithms, which have a long history in dating business cycles. Harding and Pagan (2003) also provided some comparison of univariate versions of the dating rules considered here. However, this comparison does not consider multivariate methods or the real-time performance of the methods.

In the next section we discuss the two approaches used to establish business cycle turning points in more detail. Section 3 describes the real-time dataset. Section 4 discusses the real-time performance of the models for dating turning points in the business cycle. Section 5 concludes.

2. DESCRIPTION OF THE BUSINESS CYCLE DATING METHODS

The NBER dates a turning point in the business cycle when a consensus of the Business Cycle Dating Committee that a turning point has occurred is reached. Although each Committee member likely brings different techniques to bear on this question, the decision is framed by the working definition of a business cycle provided by Arthur Burns and Wesley Mitchell (1946, p. 3):

> Business cycles are a type of fluctuation found in the aggregate economic activity of nations that organize their work mainly in business enterprises: a cycle consists of expansions occurring at about the same time in many economic activities, followed by similarly general recessions, contractions and revivals which merge into the expansion phase of the next cycle.

Fundamental to this definition is the idea that business cycles can be divided into distinct phases. In particular, expansion phases are periods when economic activity tends to trend up, whereas recession phases are periods when economic activity tends to trend down. In addition, the definition stresses that these phases are observed in many economic activities, a concept typically referred to as comovement. In practice, to date the shift from an expansion phase to a recession phase, or a business cycle peak, the NBER looks for clustering in the shifts of a broad range of series from a regime of upward trend to a regime of downward trend. The converse exercise is performed to date the shift back to an expansion phase, or a business cycle trough. Four monthly series are prominently featured by the NBER in their decisions: employment, industrial production, real manufacturing and trade sales, and real personal income excluding transfer payments.

The two business cycle dating methods that we consider in this article represent attempts to operationalize the above definition into formal algorithms and statistical models. We turn now to a more detailed discussion of both methods.

2.1 Harding and Pagan (2006) Algorithm

Based on relatively informal descriptions of NBER procedures laid out in Boehm and Moore (1984), Harding and Pagan (2006) developed a formal algorithm whereby a common set of turning points can be extracted from a group of individual time series. The algorithm is described in detail in Harding and Pagan (2006), and we provide only a brief summary here for a group of monthly time series. Before using the algorithm, we need to first extract turning point dates for each of the time series, indexed by $i = 1, \dots, I$. Here we employ the commonly used algorithm of Bry and Boschan (1971) for this purpose, which, roughly speaking, identifies turning points as local minima and maxima in the path of each time series. To implement the Bry–Boschan algorithm, we use Gauss code created for Watson (1994). Once the Bry–Boschan algorithm has been applied to each time series, we have a set of I turning point histories, labeled $\{P_1, P_2, \dots, P_I\}$ for peaks and $\{T_1, T_2, \dots, T_I\}$ for troughs, where P_i and T_i are vectors of turning point dates for time series i. The contribution of the Harding and Pagan algorithm is to consolidate these individual peak and trough dates into a single set of common turning point dates. To do this, Harding and Pagan defined variables DP_{it} and DT_{it}, which record the distance in months between month t and the nearest entry in P_i for DP_{it} and T_i for DT_{it}. For example, if $P_i = (20, 40, 60)$ and $t = 45$, then $DP_{it} = 5$. For each value of t, we then form DP_t and DT_t as the median across the I time series, that is, $DP_t = \text{median}(DP_{1t}, DP_{2t}, \dots, DP_{It})$ and $DT_t = \text{median}(DT_{1t}, DT_{2t}, \dots, DT_{It})$. Harding and Pagan then defined the common peak and trough dates as local minima in DP_t and DT_t. Formally, a common peak or trough is defined at month t if DP_t or DT_t is a minimum value in a 31-month window centered at time t, that is, from $t - 15$ to $t + 15$. In practice, these local minimum values may not be unique, and it may be necessary to break ties. To do so, Harding and Pagan considered higher percentiles than the median until a unique local minimum was found.

Finally, once the candidate set of common turning points has been obtained, two censoring procedures are applied. First, for a candidate common peak (trough) to be retained at time t, the median distance to individual turning point dates, that is, the value of DP_t (DT_t), must not be larger than 15 months. Second, turning points are recombined so that they alternate between peaks and troughs.

2.2 Dynamic Factor Markov-Switching Model

As discussed earlier, the NBER definition of a business cycle places heavy emphasis on regime shifts in economic activity. Given this, the Markov-switching model of Hamilton (1989), which endogenously estimates the timing of regime shifts in the parameters of a time series model, seems well suited for the task of modeling business cycle phase shifts. In addition, the NBER definition stresses the importance of comovement among many economic variables. This feature of the business cycle is often captured using the dynamic common factor model of Stock and Watson (1989, 1991).

Chauvet (1998) combined the dynamic-factor and Markov-switching frameworks to create a statistical model capturing both regime shifts and comovement. Specifically, defining Y_{it} as the log level of the ith time series, and $y^*_{it} = y_{it} - \bar{y}_i$ as the demeaned first difference of Y_{it}, the DFMS model has the form

$$\begin{bmatrix} y^*_{1t} \\ y^*_{2t} \\ \cdot \\ \cdot \\ y^*_{It} \end{bmatrix} = \begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \cdot \\ \cdot \\ \gamma_I \end{bmatrix} c_t + \begin{bmatrix} e_{1t} \\ e_{2t} \\ \cdot \\ \cdot \\ e_{It} \end{bmatrix}. \tag{1}$$

That is, the demeaned first difference of each series is made up of a component common to each series, given by the dynamic factor c_t, and a component idiosyncratic to each series, given by e_{it}. The common component is assumed to follow a stationary autoregressive process:

$$\phi(L)(c_t - \mu_{S_t}) = \varepsilon_t, \tag{2}$$

where ε_t is a normally distributed random variable with mean zero and variance set equal to unity for identification purposes, and $\phi(L)$ is a lag polynomial with all roots outside of the unit circle. The common component is assumed to have a switching mean, given by $\mu_{S_t} = \mu_0 + \mu_1 S_t$, where $S_t = \{0, 1\}$ is a state variable that indexes the regime and $\mu_1 < 0$ for normalization purposes. The state variable is unobserved, but is assumed to follow a Markov process with transition probabilities $P(S_t = 1|S_{t-1} = 1) = p$ and $P(S_t = 0|S_{t-1} = 0) = q$. Finally, each idiosyncratic component is assumed to follow a stationary autoregressive process:

$$\theta_i(L)e_{it} = \omega_{it}, \tag{3}$$

where $\theta_i(L)$ is a lag polynomial with all roots outside the unit circle.

Chauvet (1998) estimated the DFMS model for U.S. monthly data on nonfarm payroll employment, industrial production, real manufacturing and trade sales, and real personal income excluding transfer payments. The model produced estimated probabilities of the regime at time t conditional on the data, denoted $P(S_t = 1|\Psi_T)$, that closely matched NBER expansion and recession episodes. That is, $P(S_t = 1|\Psi_T)$ was high during recessions and low during expansions.

In this article, we use the DFMS model to obtain recession probabilities in real time. Also, because we are interested in obtaining specific turning point dates, we will require a rule to convert the recession probabilities into a zero/one variable that defines whether the economy is in an expansion or a recession regime at time t. Here, we take a conservative, two-step approach, which we outline for a business cycle peak. In the first step, we require that the probability of recession move from below to above 80% and remain above 80% for three consecutive months before a new recession phase is identified. That is, we require that $P(S_{t+k} = 1|\Psi_T) \geq .8$, for $k = 0$ to 2 and $P(S_{t-1} = 1|\Psi_T) < .8$. In the second step, the first month of this recession phase is identified as the first month prior to month t for which the probability of recession moves above 50%. That is, we find the smallest value of q for which $P(S_{t-q-1} = 1|\Psi_T) < .50$ and $P(S_{t-q} = 1|\Psi_T) \geq .50$. The peak date for this recession phase is then established as the last month of the previous expansion phase, or month $t - q - 1$. An analogous procedure, with the 80% threshold replaced by 20%, is used to establish business cycle troughs.

To estimate the parameters of the DFMS model, as well as the recession probabilities, we use the Bayesian Gibbs Sampling approach described in Kim and Nelson (1998). The Gibbs Sampler produces a posterior distribution for S_t conditional on the data, Ψ_T, the mean of which corresponds to the recession probability $P(S_t = 1|\Psi_T)$. These probabilities are then used to obtain business cycle turning point dates. Priors for the Bayesian estimation are quite diffuse, and match those used in Kim and Nelson (1998). We set the lag order of each autoregressive polynomial, $\phi(L)$ and $\theta_i(L)$, equal to 2. This choice of lag order is based on specification tests reported in the studies of Stock and Watson (1991), Chauvet (1998), and Kim and Nelson (1998), each of which suggested that 2 lags is sufficient for dynamic-factor models of the four coincident variables we consider here.

3. REAL–TIME DATASET

In this section we describe the real-time dataset. We have compiled real-time data on four coincident variables: (1) nonfarm payroll employment (EMP), (2) industrial production (IP),

(3) real manufacturing and trade sales (MTS), and (4) real personal income excluding transfer payments (PIX). These are the four monthly variables highlighted by the NBER in establishing turning point dates. We have collected realizations, or *vintages*, of these time series as they would have appeared at the end of each month from November 1976 to June 2006. For each vintage from November 1976 to January 1996, the sample collected begins in January 1959 and ends with the most recent data available for that vintage. For each vintage from February 1996 to June 2006, the sample begins in January 1967. For the series EMP, IP, and PIX, data are released for month t in month $t+1$. Thus, for these variables the sample ends in month $R-1$ for vintage R. For MTS, data are released for month t in month $t+2$. Thus, for this variable the sample ends in month $R-2$ for vintage R. We obtained the EMP and IP data series from the Federal Reserve Bank of Philadelphia real-time data archive described in Croushore and Stark (2001). Data for PIX and MTS were hand-collected as part of a larger real-time data collection project at the Federal Reserve Bank of St. Louis. This dataset is new and has not yet been used in any other applications. The Appendix provides more detail on the sources used to collect the PIX and MTS series.

4. PERFORMANCE OF THE BUSINESS CYCLE DATING METHODS

4.1 Description of Real-Time Simulation Exercise

To assess the real-time performance of the two business cycle dating methods described in Section 2, we apply these techniques to the real-time dataset described in Section 3. We assume that an analyst applies the business cycle dating methods on the final day of each month, which is soon after the release of MTS data for that monthly vintage. Thus, for each monthly vintage R, we create a monthly dataset of EMP, IP, MTS, and PIX that would have been available at the end of month R. The final month of data included in this dataset is determined by the series with the least amount of data available at vintage R. As discussed in Section 3, this final data point is month $R-2$, which is the last month for which data are available for MTS. For each vintage R, the MHP algorithm and DFMS model are applied to the dataset, and a chronology of turning point dates determined. We will be particularly interested in evidence of new turning points revealed toward the end of the sample at vintage R.

The choice to restrict the entire dataset by the series with the least data available at vintage R is a conservative assessment of the information available to the analyst. Alternatively, we could have included the month $R-1$ data for EMP, IP, and PIX in conjunction with a forecast for month $R-1$ MTS data. Although potentially fruitful, we chose not to pursue this approach here for two reasons. First of all, as will be seen later, the performance of the business cycle dating methods applied to the restricted dataset is already quite good, thus demonstrating the potential benefits of their use. Second, it is not clear that the additional information for EMP, IP, and PIX would necessarily improve the performance of the dating methods, as revisions from the first to the second release of these monthly data series, particularly EMP and IP, are often very large.

Finally, it should be noted that there are two elements of this experiment that are not "real time" in nature. First of all, whereas the parameters of the DFMS model are re-estimated at each vintage, the lag orders for the DFMS model specification remain fixed across vintages. The chosen lag orders were based on specification tests conducted in prior studies, namely, Stock and Watson (1991), Chauvet (1998), and Kim and Nelson (1998). However, because all of these studies used data not available at the earlier vintages in our dataset, for each of these earlier vintages the chosen lag orders are based on data that would not have been available at that vintage. Second, the rule used to convert recession probabilities obtained from the DFMS model into turning point dates was selected with knowledge of the estimated recession probabilities obtained using the full sample of data from the most recent vintage.

4.2 Real-Time Performance of the Business Cycle Dating Methods

We now turn to the real-time performance of the business cycle dating methods. Again, we consider vintages from November 1976 to June 2006. There are, therefore, four NBER business cycle episodes to identify in real time using these vintages, namely, the 1980, 1981–1982, 1990–1991, and 2001 recessions. We will also be interested in any "false positive" turning point dates identified by the dating methods.

Tables 1–4 describe the real-time performance of the DFMS model and the MHP algorithm. The first column gives the turning point date assigned in real time by the DFMS model or MHP algorithm. In other words, this column records the date of any new turning points established by the methods. If this turning point date has a corresponding NBER turning point, the second column gives this NBER date, and the third column records the discrepancy in months between the NBER date and the date in column 1. The fourth column gives the month in which the date in column 1 would have been available. For example, the first entry in column 4 of Table 1 is July 31, 1980. This is the first date at which the DFMS model, using the dataset available, would have revealed the January 1980 peak in column 1. The

Table 1. Business cycle peak dates obtained in real time—NBER and DFMS model

Peak date: DFMS	Peak date: NBER	Lead/lag discrepancy	Peak date available: DFMS	Peak date announced: NBER	Days ahead of NBER announcement
Jan 1980	Jan 1980	0 M	Jul 31, 1980	Jun 3, 1980	−58
Aug 1981	Jul 1981	−1 M	Feb 28, 1982	Jan 6, 1982	−53
Jul 1990	Jul 1990	0 M	Feb 28, 1991	Apr 25, 1991	56
Jan 2001	Mar 2001	2 M	Jan 31, 2002	Nov 26, 2001	−66

Table 2. Business cycle trough dates obtained in real time—NBER and DFMS model

Trough date: DFMS	Trough date: NBER	Lead/lag discrepancy	Trough date available: DFMS	Trough date announced: NBER	Days ahead of NBER announcement
Jun 1980	Jul 1980	1 M	Dec 31, 1980	Jul 8, 1981	189
Oct 1982	Nov 1982	1 M	May 31, 1983	Jul 8, 1983	38
Mar 1991	Mar 1991	0 M	Sep 30, 1991	Dec 22, 1992	449
Nov 2001	Nov 2001	0 M	Aug 31, 2002	July 17, 2003	320

fifth column gives the date the NBER announced the turning point date in column 2. The final column gives the amount of time before the NBER date that the turning point from the dating methods would have been available, which is the amount of time the date in column 4 anticipates that in column 5.

We begin with Tables 1 and 2, which show the results for the DFMS model. The DFMS model identifies eight turning points in real time, each of which corresponds to a NBER turning point. Thus, the DFMS model does not generate any false positives. The DFMS model also identifies these eight turning points with a high level of accuracy. In particular, for seven of the eight turning points, the turning point date identified in real time is within one month of the NBER date. For the remaining turning point, the peak of the 2001 recession, the date identified by the model is two months from the NBER date.

For business cycle peaks, the DFMS model does not show any systematic improvement over the NBER in the speed at which it identifies turning points. Indeed, the DFMS model would have identified the four peaks in the sample roughly one month after the NBER announcement on average, with a maximum lag time of two months. However, the DFMS model would have identified business cycle troughs much more quickly than the NBER. The average lead time for the four troughs in the sample is 249 days, or about 8 months, with a maximum lead time of 449 days for the 1991 business cycle trough. Interestingly, the increase in speed with which the DFMS algorithm identifies business cycle troughs does not come with a noticeable loss of accuracy in identifying the NBER date. Indeed, the business cycle trough dates identified in real time are all within one month of their corresponding NBER date. Given that the DFMS model treats business cycle peak and trough episodes symmetrically, its improved timeliness over the NBER for troughs but not peaks is suggestive of an asymmetry in the NBER approach. One explanation for this is that the NBER may have an asymmetric loss function for valuing errors made in establishing the dates of business cycle peaks versus troughs.

The results in Tables 1 and 2 are derived from a combination of the recession probabilities, $P(S_t = 1|\Psi_T)$, with the dating rule used to convert these recession probabilities into recession dates. For reference, Figures 1–4 plot the values of the real-time recession probabilities used to date each peak and trough in the sample. That is, these figures show a sequence of $P(S_t = 1|\Psi_T)$ that was available at the vintage for which the business cycle peak or trough was first identified.

Tables 3 and 4 report the performance of the MHP algorithm in dating turning points in real time. Similarly to the DFMS model, the MHP algorithm also identifies eight turning points, each of which corresponds to a NBER turning point date. However, these turning points are identified less accurately in general than is the case for the DFMS model. In particular, four of the turning points are at least two months from their corresponding NBER date, with the peaks of the 1980 and 2001 recessions both six months from the NBER date.

Similarly to the DFMS model, the MHP algorithm does not show any systematic improvement over the NBER in the speed with which business cycle peaks are identified, but does show an improvement in timeliness for business cycle troughs. In particular, the MHP algorithm identified the four business cycle troughs in the sample an average of 166 days, or about 5.5 months, ahead of the NBER announcement. Although still a substantial increase in timeliness, it is a smaller improvement than that achieved by the DFMS model.

4.3 Revisions of Business Cycle Dates

The NBER has made revisions to previously established business cycle turning point dates, most recently in 1975. However, the NBER's Business Cycle Dating Committee has not revised any of the eight turning point dates it has established in real time since its inception in 1978. Does this rigidity suggest that the NBER's business cycle dates are no longer consistent with the data? Or does it instead suggest that data revealed since the establishment of these turning point dates have not altered conclusions about their timing? In this section we provide some evidence on these questions.

We can evaluate the importance of data revisions for establishing business cycle turning point dates by tracking revisions

Table 3. Business cycle peak dates obtained in real time—NBER and MHP algorithm

Peak date: MHP	Peak date: NBER	Lead/lag discrepancy	Peak date available: MHP	Peak date announced: NBER	Days ahead of NBER announcement
Jul 1979	Jan 1980	6 M	May 31, 1980	Jun 3, 1980	3
May 1981	Jul 1981	2 M	Feb 28, 1982	Jan 6, 1982	−53
Jul 1990	Jul 1990	0 M	Mar 31, 1991	Apr 25, 1991	25
Sep 2000	Mar 2001	6 M	Nov 30, 2001	Nov 26, 2001	−4

Table 4. Business cycle trough dates obtained in real time—NBER and MHP algorithm

Trough date: MHP	Trough date: NBER	Lead/lag discrepancy	Trough date available: MHP	Trough date announced: NBER	Days ahead of NBER announcement
Jul 1980	Jul 1980	0 M	Apr 30, 1981	Jul 8, 1981	69
Oct 1982	Nov 1982	1 M	Jul 31, 1983	Jul 8, 1983	−23
Jul 1991	Mar 1991	−4 M	Feb 28, 1992	Dec 22, 1992	298
Oct 2001	Nov 2001	1 M	Aug 31, 2002	July 17, 2003	320

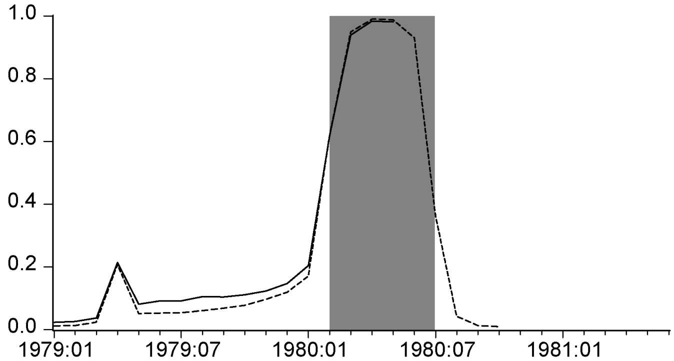

Figure 1. Real-time probabilities of recession determining the peak (——) and trough (-----) of the 1980 recession, and NBER recession (shaded).

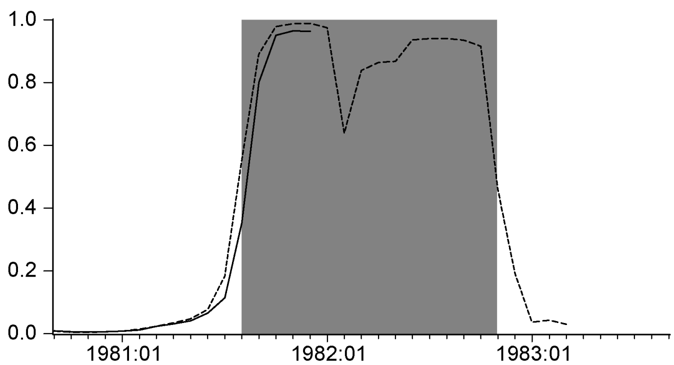

Figure 2. Real-time probabilities of recession determining the peak (——) and trough (-----) of the 1981–1982 recession, and NBER recession (shaded).

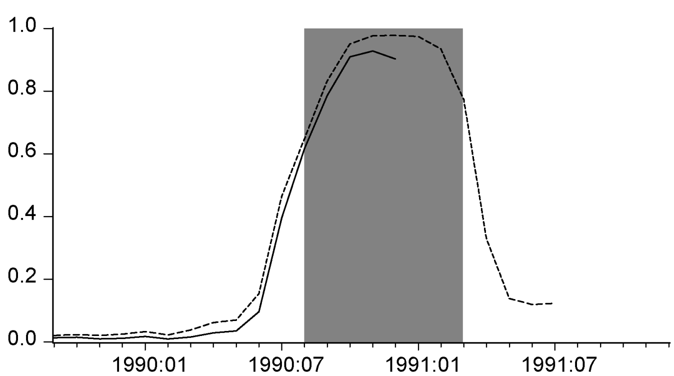

Figure 3. Real-time probabilities of recession determining the peak (——) and trough (-----) of the 1990–1991 recession, and NBER recession (shaded).

to the dates established in real time using the formal business cycle dating rules evaluated in this article. Given the superior performance of the DFMS model for mimicking the NBER dates established in real time, we focus on this approach. In particular, we apply the DFMS model to the most recent vintage of data available in our dataset, June 2006, and obtain a chronology of business cycle turning point dates. We then compare the business cycle turning point dates established in real time by the DFMS model to those established using the most recent vintage of data. Table 5 contains this comparison.

The results in Table 5 demonstrate that in most cases, data revisions do not appear to be an important factor for determining the timing of business cycle turning points. In particular, for seven of the eight turning points in the sample, the date established by the DFMS model using the final vintage of data available is within one month of that established in real time. Indeed, for four of the eight turning points there is no revision to the turning point date established in real time.

The single case where the real-time business cycle date is revised by more than one month, namely, the peak of the 2001 recession, merits further discussion. Note that the peak of the 2001 recession is established by the DFMS model in real time to be January of 2001, two months prior to the March 2001 peak established by the NBER. From Table 1, this peak date would not have been available from the DFMS model until two months after the official announcement by the NBER. Thus, the initial date established by the DFMS model is already based on more information than was available to the NBER. Further, this peak date is moved an additional two months earlier, to November of 2000, when the DFMS model is applied to the June 2006 vintage of data. Note that data available in June 2006 is not necessary for the DFMS model to make this revision. In particular,

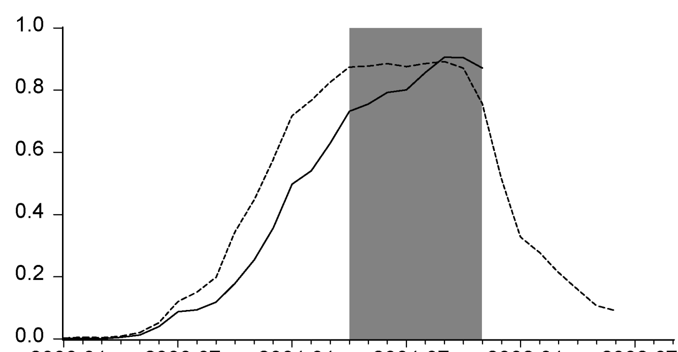

Figure 4. Real-time probabilities of recession determining the peak (——) and trough (-----) of the 2001 recession, and NBER recession (shaded).

Table 5. Revisions to business cycle dates: DFMS model

Peaks			Troughs		
NBER date	Initial date: DFMS	Final date: DFMS	NBER date	Initial date: DFMS	Final date: DFMS
Jan 1980	Jan 1980	Jan 1980	Jul 1980	Jun 1980	Jun 1980
Jul 1981	Aug 1981	Jul 1981	Nov 1982	Oct 1982	Nov 1982
Jul 1990	Jul 1990	Aug 1990	Mar 1991	Mar 1991	Mar 1991
Mar 2001	Jan 2001	Nov 2000	Nov 2001	Nov 2001	Nov 2001

the revision to November of 2000 would have first been available from the DFMS model by the July 2002 vintage. In sum, data revealed after the official announcement by the NBER of the March 2001 peak seem to be consistent with this peak occurring somewhat earlier, and provide one example suggestive that an established NBER date may be inconsistent with revised data.

Although not revealed in Table 5, the trough of the 2001 recession is also an interesting case for investigating the effects of additional and revised data on conclusions about turning point dates. In particular, from Table 2, the DFMS model would have first established the trough date of November 2001 by the end of August of 2002. However, for a brief period for vintages in mid-2003, the recession probabilities from the DFMS model for 2002 and 2003 rose significantly to levels consistent with a continuation of the 2001 recession. This was the result of very weak employment data observed in 2002 and 2003, or the so-called "jobless recovery." By the end of 2003, the recession probabilities would have returned to levels consistent with the previously established trough date of November 2001. This episode demonstrates that the caution exercised by the NBER in establishing the trough of the 2001 recession may have been justified, particularly if their primary objective is to establish turning point dates that are unlikely to need revision.

5. CONCLUSIONS

This article investigates the ability of formal rules to establish business cycle turning point dates in real time. Both methods studied, a nonparametric algorithm given in Harding and Pagan (2003) and the dynamic-factor Markov-switching model as in Chauvet (1998), identify the NBER turning point dates in real time with reasonable accuracy, and with no instances of false positives. Both approaches also provide improvements over the NBER in the timeliness with which they identify business cycle troughs, but provide no such improvement for business cycle peaks. Comparing the two methods, the dynamic-factor Markov-switching model identifies NBER turning point dates more accurately, as well as identifies business cycle troughs with a larger lead.

ACKNOWLEDGMENTS

We thank Jon Faust, Robert Rasche, two anonymous referees, and seminar participants at the 2005 Winter Meetings of the Econometric Society for helpful comments. Research assistance by Michelle Armesto and Garrett Holland was invaluable in the completion of this project. We owe special thanks to Robert Rasche for his assistance in obtaining the real-time dataset. We are also grateful to Don Harding for sharing his Gauss code. Much of this article was completed while Piger was a Senior Economist at the Federal Reserve Bank of St. Louis. The views expressed in this article should not be interpreted as those of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

APPENDIX: SOURCES OF REAL–TIME DATA

A.1 Real Personal Income Excluding Transfer Payments

For vintages from November 1976 through March 1990, data for real personal income excluding transfer payments were collected from *Business Conditions Digest.* For vintages from April 1990 through December 1995, data for real personal income excluding transfer payments were collected from the *Survey of Current Business*. For vintages from January 1996 through June 2006, nominal personal income, nominal disposable personal income, and real disposable personal income were collected from the Federal Reserve Bank of St. Louis ALFRED database, whereas data for nominal transfer payments were collected from *Economic Indicators*, *Business Statistics*, the *Survey of Current Business*, and data archives maintained by the Federal Reserve Bank of St. Louis. Data for real personal income excluding transfer payments were then formed by subtracting nominal transfer payments from nominal personal income, and dividing by the ratio of nominal to real disposable personal income.

A.2 Real Manufacturing and Trade Sales

For vintages from November 1976 through March 1990, data for real manufacturing and trade sales were collected from *Business Conditions Digest*, whereas for vintages from April 1990 through December 1995, real manufacturing and trade sales data were collected from the *Survey of Current Business*. For vintages from January 1996 through June 2006, real manufacturing and trade sales data were collected from *Business Cycle Indicators*, *Business Statistics*, the *Survey of Current Business*, and data archives maintained by the Federal Reserve Bank of St. Louis.

For a small number of individual vintages, there were gaps in the data available. These missing data were filled in using the following strategy. Suppose that for the Rth vintage, we are missing data from period t to $t+k$. Denote these missing data as $Y_t^R, Y_{t+1}^R, \ldots, Y_{t+k}^R$. Suppose that data are available for

$Y_{t-1}^{R-h}, Y_t^{R-h}, \ldots, Y_{t+k}^{R-h}, Y_{t+k+1}^{R-h}$, as well as for Y_{t-1}^R and Y_{t+k+1}^R. Our imputed value for Y_j^R, denoted $\hat{Y}_j^R$, is then given by

$$\hat{Y}_j^R = \hat{Y}_{j-1}^R \frac{Y_j^{R-h}}{Y_{j-1}^{R-h}} \left(\frac{r_1}{r_2}\right)^{1/(k+2)}, \qquad j = t, \ldots, t+k,$$

where $r_1 = Y_{t+k+1}^R / Y_{t-1}^R$, $r_2 = Y_{t+k+1}^{R-h} / Y_{t-1}^{R-h}$, and the recursion is initialized with $\hat{Y}_{t-1}^R = Y_{t-1}^R$.

In words, this imputation formula fills in the missing data for period j using the actual growth rate observed in period j from the data recorded at vintage $R-h$ (the first bracketed term) modified by an amount that does not vary with j (the second bracketed term). This modification ensures that the difference in total growth observed from period $t-1$ to period $t+k+1$ using data from vintages R and $R-h$ is spread evenly over the period t to $t+k+1$.

[Received April 2005. Revised February 2007.]

REFERENCES

Boehm, E., and Moore, G. H. (1984), "New Economic Indicators for Australia, 1949–1984," *Australian Economic Review*, 68, 34–56.

Bry, G., and Boschan, C. (1971), *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs*, New York: National Bureau of Economic Research.

Burns, A. F., and Mitchell, W. A. (1946), *Measuring Business Cycles*, New York: National Bureau of Economic Research.

Chauvet, M. (1998), "An Econometric Characterization of Business Cycle Dynamics With Factor Structure and Regime Switching," *International Economic Review*, 39, 969–996.

Chauvet, M., and Hamilton, J. (2006), "Dating Business Cycle Turning Points in Real Time," in *Nonlinear Time Series Analysis of Business Cycles*, eds. C. Milas, P. Rothman, and D. Van Dijk, Amsterdam: Elsevier Science, pp. 1–54.

Chauvet, M., and Piger, J. (2003), "Identifying Business Cycle Turning Points in Real Time," *Federal Reserve Bank of St. Louis Review*, 85, 47–61.

Croushore, D., and Stark, T. (2001), "A Real-Time Data Set for Macroeconomists," *Journal of Econometrics*, 105, 111–130.

Hamilton, J. D. (1989), "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357–384.

Harding, D., and Pagan, A. (2003), "A Comparison of Two Business Cycle Dating Methods," *Journal of Economic Dynamics and Control*, 27, 1681–1690.

——— (2006), "Synchronization of Cycles," *Journal of Econometrics*, 132, 59–79.

Kim, C.-J., and Nelson, C. R. (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.

Layton, A. P. (1996), "Dating and Predicting Phase Changes in the U.S. Business Cycle," *International Journal of Forecasting*, 12, 417–428.

Stock, J. H., and Watson, M. W. (1989), "New Indexes of Coincident and Leading Economic Indicators," *NBER Macroeconomics Annual*, 4, 351–393.

——— (1991), "A Probability Model of the Coincident Economic Indicators," in *Leading Economic Indicators: New Approaches and Forecasting Records*, eds. K. Lahiri and G. H. Moore, Cambridge, U.K.: Cambridge University Press, pp. 63–90.

Watson, M. W. (1994), "Business Cycle Durations and Postwar Stabilization of the U.S. Economy," *American Economic Review*, 84, 24–46.