E. Andersson*, D. Bock, M. Frisén

Some statistical aspects of methods for detection of turning points in business cycles

Journal of Applied Statistics, Volume 33, Issue 3 April 2006 , pages 257 - 278

http://dx.doi.org/10.1080/02664760500445517

Some Statistical Aspects of Methods for Detection of Turning Points in Business Cycles

E. Andersson*, D. Bock, M. Frisén, *Statistical Research Unit, Goteborg University, Sweden*

Abstract *Methods for on-line turning point detection in business cycles are discussed. The statistical properties of three likelihood based methods are compared. One is based on a Hidden Markov Model, another includes a non-parametric estimation procedure and the third combines features of the other two. The methods are illustrated by monitoring a period of the Swedish industrial production. Evaluation measures that reflect timeliness are used. The effects of smoothing, seasonal variation, autoregression and multivariate issues on methods for timely detection are discussed.*

Key words: *Monitoring; Surveillance; Early warning system; Regime switching; Non-parametric*.

* Corresponding author is E. Andersson, Statistical Research Unit, PO Box 660, SE-405 30 Goteborg, Sweden.

1. Introduction

For both government and industry it is important to have systems for predicting the future state of the economy, for example timely prediction of a change from a period of expansion to one of recession. Leading economic indicators can be used to predict the turns of the business cycles. By constructing a system for early warnings about turns in one or several leading indicators, the turning point time of the general business cycle can be predicted.

Here, we study different methods for early warning, i.e. methods for timely detection of a regime shift, from a recession phase to an expansion phase (or vice versa) in a leading index. For reviews and general discussions on the importance of timeliness in the detection of turning points see e.g. Neftci (1982), Zarnowitz and Moore (1982), Westlund and Zackrisson (1986), Hackl and Westlund (1989), Zellner et al. (1991), Li and Dorfman (1996) and Layton and Katsuura (2001). Suggested methods for detecting regime shifts are presented in the works by e.g. Diebold and Rudebusch (1989), Hamilton (1989), Jun and Joo (1993), Lahiri and Wang (1994), Layton (1996), Koskinen and Öller (2003), Layton (1998) and Birchenhall et al. (1999).

As pointed out by e.g. Diebold and Rudebusch (1996), Kim and Nelson (1998) and Birchenhall et al. (1999) two distinct but related approaches to the characterization and dating of the business cycle can be found. One approach emphasizes the common movements of several variables. This approach is pursued by e.g. Stock and Watson (1991) and Stock and Watson (1993) and is briefly discussed in Section 3.2.5 on multivariate approaches. The other approach, the regime shift, is the one pursued in this paper, as also in the works by Neftci (1982), Diebold and Rudebusch (1989), Hamilton (1989), Jun and Joo (1993), Lahiri and Wang (1994), Layton (1996), Koskinen and Öller (2003), Layton (1998) and Birchenhall et al. (1999).

The inference situation in this paper is one of repeated decisions: a new decision is made after each new observation. Several other research areas treat turning points, but from another perspective, for example estimating the location and size of structural breaks in a series of fixed number of observations (Mudambi (1997), Delgado and Hidalgo (2000)). To make this decision an alarm system is used, based on an alarm statistic and an alarm limit. Statistical surveillance, and especially the inference questions regarding the repeated decisions, has been investigated and developed by e.g. Shiryaev (1963), Frisén and de Maré (1991), Wetherhill and Brown (1991), Srivastava and Wu (1993), Lai (1995), Frisén and Wessman (1999) and Frisén (2003).

Statistical surveillance methods were first used in industrial process control, but are now used in many areas, for example for financial decisions (signaling the optimal time to trade). In Theodossiou (1993) and Blondell et al. (2002) a surveillance method, CUSUM, is used to discriminate between healthy firms and firms in financial distress. Since September 11, 2001, there is an increased interest in surveillance methods, in order to detect bio terrorism (on-line detection of health hazards).

In recent years on-line surveillance methods based on the likelihood ratio (or posterior distribution) have been in focus. In the general theory on statistical surveillance there are proofs for their optimality properties (see e.g. Shiryaev (1963) and Frisén and de Maré (1991)). The posterior probability is often used as an alarm statistic in methods based on a hidden Markov model (HMM), see Hamilton (1989).

When comparing and evaluating methods for on-line detection, it is important to consider timeliness, i.e. the time of the alarm in relation to the time of the turn. Suitable measures are the probability of successful detection within a specified time limit, and the expected delay of an alarm. Another important aspect is the predictive value of an alarm, which can indicate the appropriate action

One purpose of this paper is to review some likelihood ratio based surveillance methods. We discuss different assumptions regarding the process under surveillance and the estimation procedures connected with them. Data transformations, e.g. seasonal adjustment and smoothing can distort the characteristics of the data and hence influence the surveillance performance. Some of the suggested approaches for dealing with seasonal variation are reviewed and the effects of smoothing the data are investigated by a Monte Carlo study. The purpose is also to discuss how to deal with autocorrelated data, trend and multivariate problems in surveillance. The properties of the methods are demonstrated by application to the problem of turning point detection in the Swedish Industrial Production (IP).

The paper is organized as follows. Section 2 contains a description of different likelihood based approaches in surveillance and common ways in surveillance for making different methods comparable. In Section 3 special data problems and estimation procedures are discussed. Also a Monte Carlo study regarding the effects of smoothing on the detection ability is presented in this section. In Section 4, three monitoring methods are applied to a period of the Swedish industrial production containing a turn and the pros and cons of this way to evaluate methods are discussed. Also in this section, the timeliness properties of the three methods are described. Section 5 contains a summarizing discussion.

2. Concepts of likelihood based surveillance for detection of turning points

In this section the basic concepts of likelihood based surveillance for online detection of turning points are given.

First, we formulate the surveillance problem. Here, a process X (leading economic indicator) is under surveillance, where X is often measured monthly or quarterly. For quarterly data, a new observation becomes available every quarter. Thus every quarter, based on the available observations, we decide whether the observations so far indicate a turn. An alarm system is developed for this purpose, with an alarm statistic and an alarm limit. The alarm statistic at time s is a function of the available observations, $p(x_s) = p(x(1), ..., x(s))$ and the alarm limit can be constant or time dependent. In the alarm statistic the observations $x(1), ..., x(s)$ are weighted together. For example in the EWMA method of surveillance (Roberts (1959)) the observations are weighted exponentially, whereas in the Shewhart method (Shewhart (1931)), all weight is given to the last observation. The EWMA method is evaluated in e.g. Sonesson (2003).

Thus at every decision time s, we use the alarm system to decide whether there has been a turn or not. This can be formulated as, at every decision time, discriminating between two events: D ="the turn has not occurred yet" and C ="the turn has occurred". This is further described in the next section.

2.1 Event to be detected

The situation under study is one where X is a leading economic indicator. By monitoring X we want to detect a regime shift (a turning point) as soon as possible. The model for X at time t is:

$$X(t) = \mu(t) + \varepsilon(t), \quad (1)$$

where $\varepsilon(t)$ ~iid N[0; σ^2] and $\mu(t)$ is a cyclical stochastic process described below. This simple assumption regarding the disturbance term is used to emphasize the inferential issues. Suggestions of how to adjust the surveillance system in the presence of autoregression in the disturbance term are discussed in Section 3.2.3.

The regime shift occurs at an unknown and random time τ, i.e. at τ there is a turn in μ. An alarm system is developed, for on-line detection of the turn. At each decision time s an alarm statistic is used to discriminate between the two events *D(s)* ="the turn has not occurred yet" and *C(s)* ="the turn has occurred". The event *D* is often defined as $D(s) = \{\tau > s\}$. Regarding the event *C(s)*, it could be that we want to decide if a turn has occurred at a the time point *t'*, and then $C(s)=\{\tau = t'\}$. But when we want to decide whether there has been a turn since the start of the surveillance, then $C(s) = \{\tau \leq s\}$.

Knowledge of whether the next turn will be a trough or a peak is assumed. The solutions for peak- and trough-detection are equivalent, as everything is symmetrical. It is the knowledge per se which is important. For simplicity in the presentations henceforward the turning point will be expressed as a peak (a transition from expansion to recession). This is not, however, a restriction in the methods.

A time τ there is a turn in μ and as τ is random, so is μ. Different assumptions can be made about μ, conditional on *D* and *C*., i.e. about the expected value of the process in expansion (or recession) and at a turn. Parametric assumptions regarding μ make the method more powerful if the assumptions are valid. Under the assumption that the regression consists of linear functions where the slopes are symmetrical for the two phases, the aim is to discriminate between *D* and *C*, such that

$$D(s):\ \mu(s) = \beta_0 + \beta_1 \cdot s \qquad (2)$$

$$C(s) = \{\cup\, C\tau\},$$

where $C\tau:\ \mu(s) = \beta_0 + \beta_1\cdot(\tau\text{-}1) - \beta_1\cdot(s\text{-}\tau+1)$ and where $\tau=\{1, 2, ..., s\}$ and β_0 and β_1 are known constants.

Then we have the following model for the turn

$$\mathrm{E}[X(t)]=\mu(t) = \begin{cases} \beta_0 + \beta_1 \cdot t, & t < \tau \\ \beta_0 + \beta_1 \cdot (\tau - 1) - \beta_2 \cdot (t - \tau + 1), & t \geq \tau \end{cases}, \qquad (3)$$

where $t=\{1, 2, ...\}$. The expected value in (3) holds for a random walk with drift where the value of the drift parameter changes from β_1 to $-\beta_2$ at time $t=\tau$.

The assumption that μ is known is not always realistic in on-line detection. Instead of assuming that the parametric shape is known, we can use only monotonicity restrictions to define μ under *C* and *D*. Then the aim is to discriminate between the following two events:

$$D(s):\ \mu(1) \leq ... \leq \mu(s) \qquad (4)$$
$$C(s):\ \mu(1) \leq ... \leq \mu(\tau\text{-}1) \text{ and } \mu(\tau\text{-}1) \geq \mu(\tau) \geq ... \geq \mu(s)$$

where $\tau=\{1, 2, ..., s\}$ and at least one inequality is strict in the second part.

Thus here the exact parametric shape of μ is unknown. We only know that μ is monotonic within each phase, that is

$$\mathrm{E}[X_t]=\mu_t:\ \begin{cases} \mu(1) \leq \mu(2) \leq ... \leq \mu(t), & t < \tau \\ \mu(1) \leq ... \leq \mu(\tau - 1) \ and\ \mu(\tau - 1) \geq ... \geq \mu(t), & t \geq \tau \end{cases} \qquad (5)$$

where t=1 is in a period of expansion, τ is the random time of a turning point (the time of change from the expansion to a recession) and $X_t=\{X(1), X(2), ..., X(t)\}$ and where, for $t \geq \tau$, at least one inequality is strict in the second part. The monotonicity restrictions for a trough are the opposite of those in (5).

If an HMM is assumed then, at decision time s, an alarm statistic is used to discriminate between

$$D(s): \mu(s\text{-}1) \le \mu(s), \qquad (6)$$
$$C(s): \mu(s\text{-}1) > \mu(s).$$

The s in (2), (4) and (6) is the decision time and s=1 when the surveillance is started. At the next time point (e.g. the next quarter), s=2. Thus the events D and C are not constant from one decision time to the next. The difference between D and C in (2) and (4) is only the assumptions regarding $\mu(t)$. However, when D and C are specified as in (6), the events are different also in another aspect. The apparently simpler event in the HMM approach is combined with a more complicated situation for the information of previous states. No knowledge of previous states is utilized in an HMM approach. Thus the probabilities for the history of those earlier states will have an effect. The two events in (6) correspond to families of histories of states, for example are the events in (4) only a subgroup of the events in (6). Because of Markov dependence the probabilities for the histories of those earlier states will have an effect and earlier observations carry information of the history of states.

In many HMM approaches the series under observation, X, is differentiated and the expected value of the differentiated process is assumed to be constant, conditional on the state, see e.g. Layton (1996), Ivanova et al. (2000) and Layton and Katsuura (2001). If the process is assumed to be a random walk with drift, or another process with an expected value as in (3), then the differentiated series will have a constant expected value, conditional on the state. For the situation when the turn is a peak we have

$$\mathrm{E}[X(t)\text{-}X(t\text{-}1)] = \begin{cases} \beta_1, & t < \tau \\ \beta_2, & t \ge \tau \end{cases}, \qquad (7)$$

where $\beta_1 \ge 0$ and $\beta_2 < 0$. When the observations are independent over time, as in (1), the expected values in (7) imply the linear functions in (3). For a process with an expected value as in (3), then (7) is valid.

The variation might be different for recession and expansion, and French and Sichel (1993) find that the variation is largest around business-cycle troughs. Macroeconomic time series are sometimes considered to have a continuously varying standard deviation. If there is evidence of considerable heteroscedasticity, then the observations in the alarm statistic should have different weights. Another suggestion, by Fang and Zhang (1999), is to use time varying limits. Maravelakis et al. (2004) show that ARL[1] changes with a larger heteroscedasticity. Sometimes the logarithm transformation is used for variance stabilization. This is the case here, where the observation X is the logarithm of the original observation. After this transformation, the variance is here assumed equal, as also by Andersson (2004). The surveillance is conducted and evaluated for the transformed variable X.

In this paper, the alarm statistics are based on the likelihood ratio between the events D and C, i.e.

$$\frac{f(x_s|C)}{f(x_s|D)},$$

where $x_s=\{x(1), \ldots, x(s)\}$ and f is the likelihood function. This is described further in Section 2.3.

2.2 Assumptions about transition probabilities

The probability of a transition from recession to expansion (or vice versa) are, in most approaches in the HMM framework, assumed to be constant with respect to time, see e.g. Hamilton (1989), Layton (1996) and Ivanova et al. (2000). But the transition probability can also be assumed to be time varying, as by Neftci (1982), Diebold et al. (1994), Filardo (1994) and Layton and Katsuura (2001). One approach for the time varying transition probability is to use a model where the transition probability is a function of economic indicators, see e.g. Layton and Katsuura (2001). The probability of a transition might depend on the time spent in the current state and Zuehlke (2003) finds evidence of duration dependence in several US economic indicators. Filardo and Gordon (1998) use an extension of Markov switching models where a probit model allows the transition probabilities to vary with the information in the leading indicators. The assumption of a time invariant transition probability is made for all three methods investigated in this paper.

The specification of *C* and *D* in (2) and (4) implies that we know the type of the latest turn (whether it was a peak or a trough) and we specify *C* and *D* accordingly, that is if the latest turn was a trough, we want to detect the next peak and we specify *C* and *D* accordingly. At an alarm the surveillance is restarted so that we are always looking for the next turn, no more than one turn at a time. When the information about the time and type (peak or trough) of the last turning point is not utilized (as in (6) above), it is necessary to make a probability statement regarding the type of the next turning point as well as inference about whether the turning point has occurred or not. For this purpose, it is necessary to consider all previous possible turns (both peaks and troughs) and hence two transitions probabilities are needed in the monitoring system. The probabilities of transitions can also be expressed as intensities of the occurrences of peaks and troughs.

Contrary to the situation of an HMM, where a whole history is considered, sometimes the specific type of the next turn (peak or trough) is sometimes used (e.g. Neftci (1982)). In that case it is sufficient with one measure of intensity in the monitoring system. The intensity, is hereafter denoted ν.

$$\nu = P(C(t)|D(t-1)) = P\left(\tau = t|\tau \geq t\right). \quad (8)$$

Note that the value of the intensity may differ for peak detection and trough detection. The assumption of a constant transition probability, and thus a geometric distribution for the turning point time τ, is not very realistic in the business cycle application. The lengths of the cycles vary more than for a geometrical distribution and the probability of small values of the time for the turning point is much smaller for the business cycle than for a geometrically distributed variable, as shown in Figure 1.

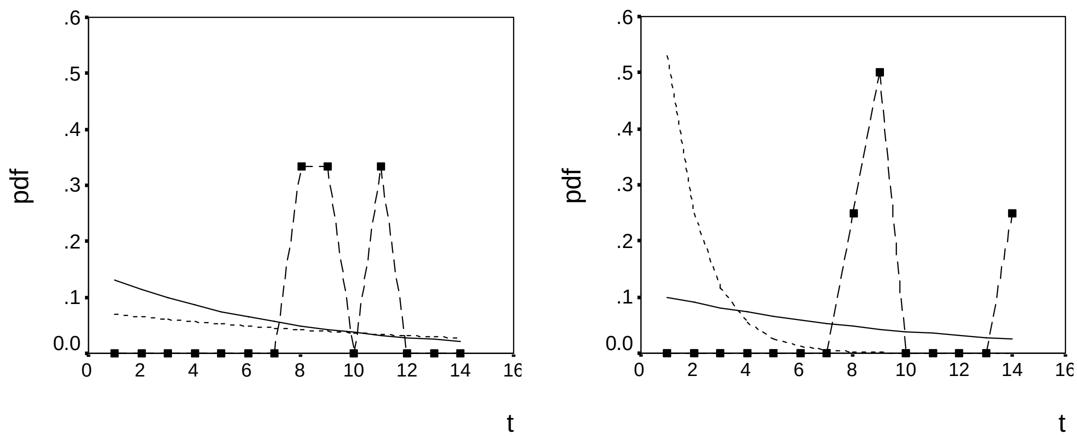

Figure 1. Left: density for peak with observed sample density function (— — — ■) compared to geometric density function with intensities 0.07 (· · ·) and 0.13 (——). Right: density for trough with observed data (— — — ■) compared to geometric density with intensities 0.53 (· · ·) and 0.10 (——). Source: National institute of Economic Research, Sweden.

It could be argued that the empirical distribution for τ should be used. However a prior based on the observed data above would result in a high prior probability for a turning point after about 9 quarters. The consequence would be that the influence of the actual data is reduced and the probability of an alarm after 9 quarters would be very high, only due to prior information.

Estimation of the probabilities can be made using data from an earlier period. But a very long time series is required since there is often several years between transitions in the business cycle (see Figure 1).

In order to avoid the risk of misspecification, a non-informative prior for the turning point time can be used. This is done by the Shiryaev-Roberts (SR) approach, suggested by Shiryaev (1963) and Roberts (1966). The SR approach is used for two of the methods in this paper.

2.3 Alarm rules

As mentioned in Section 2.1, all methods considered here use a likelihood ratio based alarm statistic. The likelihood ratio (LR) method has several optimal properties, see Frisén and de Maré (1991). The expected utility, based on very general functions of the gain of an alarm and the loss of a false alarm, is maximized. The LR method yields a minimum expected delay of an alarm signal conditional on a fixed probability of false alarm. In Frisén and Wessman (1999) several properties of the LR method are investigated and compared with other methods of surveillance, e.g. the Shewhart method and the CUSUM method (Page (1954)). Frisén and de Maré (1991) showed that the posterior probability approach is equivalent to the LR approach for the situation where *C* is the complement of *D*.

The LR method is based on the full likelihood which is a weighted sum of the partial likelihoods for the components in *C*={*C1, C2, ..., Cs*}, For the situation where μ are known functions under *D* and *C* (see (3)), the alarm rule for the likelihood ratio at time s is written as

$$\sum_{j=1}^{s} \frac{w_s(j)\cdot f(x_s \,|\, \mu=\mu^{Cj})}{f(x_s \,|\, \mu=\mu^{D})} > k_s,$$

where $w_s(j)$ = P(τ =j)/P($\tau \leq s$), k_s=k/(1-k)·P(τ >s)/P($\tau \leq s$) and μ^{Cj} and μ^D are vectors under restriction Cj={τ=j} and D={τ>s}.

The assumptions made about the intensity affect the alarm method through the weights, $w_s(j)$, and the alarm limits, k_s. In the SR approach, the weights as well as the alarm limit are constant.

$$\sum_{j=1}^{s} \frac{f(x_s | \mu = \mu^{Cj})}{f(x_s | \mu = \mu^{D})} > k_{SRlin}, \tag{9}$$

where k_{SRlin} is a constant alarm limit. The method, where μ^D and μ^{Cj} are modelled using known linear functions with a symmetric turning point and where the SR approach is used regarding the intensity, is hereafter referred to as the *SRlin* method. The SR approach (i.e. with equal weights) has also been used for a Poisson process (Kenett and Pollak (1996)).

The non-parametric approach, without parametric assumptions regarding μ, has the alarm rule

$$\sum_{j=1}^{s} \frac{f(x_s | \mu = \hat{\mu}^{Cj})}{f(x_s | \mu = \hat{\mu}^{D})} > k_{SRnp}, \tag{10}$$

where k_{SRnp} is a constant alarm limit and $\hat{\mu}^D$ and $\hat{\mu}^{Cj}$ are the maximum likelihood estimators of the vector μ under monotonicity restrictions D and Cj, described in Section 3.1.1. This method is hereafter referred to as the *SRnp* method and was suggested by Frisén (1994) and evaluated by Andersson (2002) and Andersson et al. (2004).

With an HMM approach (e.g. Hamilton 1989), the posterior probability is used in order to classify time points into either expansion or recession, based on the following rule

$$P(C(s)|x_s) > k_{HMM}. \tag{11}$$

The alarm limit, k_{HMM}, is usually chosen to be 0.5 (see e.g. Hamilton (1989) and Ivanova et al. (2000)). The classification by the posterior probability can be used in prospective monitoring (see e.g. Neftci (1982)). LeSage (1991) uses the posterior probability as a turning-point indicator. By "the *HMlin* method" we hereafter refer to a monitoring method where the rule in (11) is used with k_{HMM}=0.5, together with the definition of the event C in (6) and the assumption that the differentiated series is constant in each regime according to (7). These conditions agree with those used by Koskinen and Öller (2003).

The approach by Birchenhall et al. (1999) is similar to both the HMM approach and the likelihood ratio method of surveillance in two respects: i) Birchenhall's approach is based on Bayes theorem and the likelihood and ii) a classification is made of the type of regime. A major difference, however, is that the classification into different regimes is based on explaining variables and not on the earlier state. This difference is discussed further in Andersson et al. (2004).

The time of alarm, t_A, is the first time for which the alarm statistic (in (9), (10) and (11) respectively) exceeds its specified alarm limit.

The alarm limits can be determined indirectly, in order to control the false alarms. The most common way in the general theory and practice of surveillance is to control the ARL^0, (the Average Run Length to the first alarm if the process has no turn). Also the MRL^0 (the median run length) has been used (Hawkins (1992), Gan (1993) and Andersson (2002)), which has a clearer interpretation for skewed distributions. The alarm limit 0.5, for the posterior probability (11), is based on a symmetric loss function and no direct conclusion can be made regarding the rate of false alarms. Canova and Ciccarelli (2004) provide methods for forecasting variables and turning points using VAR models and an alarm is called if the probability of a turn, given a model, exceeds 0.5.

3. Estimation and special data problems

When a surveillance method is applied to a set of data, for example the industrial production, in order to detect the next turn, the user is faced with several practical problems: Which assumption should be made regarding the parameters of the surveillance system? Which estimation procedure should be used? If the data is very noisy – should it be smoothed before monitoring? How will seasonal adjustment and trend adjustment affect the performance of the surveillance? How should surveillance be made if we have more than one leading economic indicator? For all these questions the special situation of surveillance must be born in mind – we can only use data from previous cycles for estimation, since the current cycle is not over yet.

3.1 Estimation

An important aspect is which assumptions that are made about the process, since these assumptions determine how the parameters of the model are estimated. The parameters are often estimated using previous data. If the parameters are estimated from a short period the variance of the estimates will be large and the parameters might be severely misspecified. As a result the method will produce misguiding results, leading to wrong conclusions and decisions. Here we discuss assumptions and estimation of μ and the transition probabilities.

3.1.1 Estimation of the level

Sometimes it is assumed that μ is known, which, in practice, means that it has been estimated from a large enough set of data from earlier periods.

If it is assumed that the differentiated series has a constant expected value, conditional on the state, and that the expected values are constant over the cycles, then the estimation can be made using previous data. This assumption is used by e.g. Neftci (1982), Layton (1996), Ivanova et al. (2000) and Layton and Katsuura (2001). One example of such an estimation procedure, under the model assumptions in (1), is to first use some classification rule to classify each time point in the estimation period as belonging to either the expansion state or the recession state (see e.g. Koskinen and Öller (2003)). Then the parameters β_1 and β_2 in (7) can be estimated as

$$\hat{\beta}_j = \bar{d}_j = \frac{1}{n_j}\sum_{i=1}^{nj} d_{ji} ,$$

where n_j = # time points classified as state j and d_j is the differentiated series classified as belonging to state j.

In order to avoid the rather strong assumption of a specific parametric function for μ, or when reliable information on the parametric function is not available, an approach based only on the knowledge that the monotonicity of μ changes at a turning point can be used, as specified in (5). Then μ is also estimated, however not under a strong parametric assumption, but only under the monotonicity restrictions in (5). The estimation is made using a least square criterion and under the model assumptions in (1), the estimates $\hat{\mu}^D$ and $\hat{\mu}^C$ are also the maximum likelihood estimates. The $\hat{\mu}^D$ is the estimated parameter vector which corresponds to

$$\max_{\mu \in F^D} f(x_s | \mu),$$

where F^D is the family of μ-vectors such that $\{\mu(1) \leq \mu(2) \leq \ldots \leq \mu(s)\}$. Thus, $\hat{\mu}^D$ is the maximum likelihood estimator of μ under the monotonicity restriction *D*. For a trough, the estimation is made under the restriction $\{\mu(1) \geq \ldots \geq \mu(s)\}$. This estimator is described by e.g. Robertson et al. (1988), p. 1-58.

The event *C* is composite, $C=\{\tau \leq s\}$, and thus we have $C = \{C1, C2, \ldots, Cs\}$ and $\hat{\mu}^{Cj}$ is the estimated parameter vector which corresponds to

$$\max_{\mu \in F^{Cj}} f(x_s | \mu),$$

where F^{Cj} is the family of μ -vectors such that $\{\mu(1) \leq \ldots \leq \mu(j\text{-}1) \text{ and } \mu(j\text{-}1) \geq \mu(j) \geq \ldots\}$, where $j = \{1, 2, \ldots, s\}$ and where at least one inequality is strict in the second part. Thus, $\hat{\mu}^{Cj}$, $j \in \{1, 2, \ldots, s\}$, is the maximum likelihood estimator of μ under the monotonicity restriction *Cj*. This estimator is given by Frisén (1986). For a trough, the estimation is made under the restriction $\{\mu(1) \geq \ldots \geq \mu(j\text{-}1) \text{ and } \mu(j\text{-}1) \leq \mu(j) \leq \ldots\}$.

3.1.2 Estimation of transition probabilities

When considering different methods of estimation, simultaneous maximum likelihood estimation of all parameters in the model is an obvious choice. However, if the whole parameter set is estimated using a maximum likelihood criterion then the rareness of the turning points can lead to large errors around turning points, compensated by high accuracy within phases, as pointed out by Lahiri and Wang (1994) and Koskinen and Öller (2003). For that reason, the transition probabilities are sometimes estimated using some other criterion than maximum overall likelihood.

Of the three surveillance methods compared here, only the HMlin method needs estimates of the transition probabilities. Maximum likelihood estimates, based only on the events of transitions, are a natural choice. For the data on the Swedish IP (source: National Institute of Economic Research, Sweden), the transition probabilities, p_{12} and p_{21}, are estimated using

$$\hat{p}_{12} = \frac{n_{12}}{n_{11} + n_{12}} = \frac{5}{34+5} = 0.13 \ (0.054), \qquad (12)$$

$$\hat{p}_{21} = \frac{n_{21}}{n_{21} + n_{22}} = \frac{4}{4+36} = 0.10 \ (0.047),$$

where n_{ij} is the number of transitions from state *i* to state *j* and the standard errors are given in parenthesis.

3.2 Special data problems

There are many special data problems when applying methods for on-line turning point detection in cyclical, economic processes. In this section we discuss the problems connected with using a surveillance system on data that have been transformed (e.g. smoothed, adjusted for seasonality and adjusted for trend). The section is also concerned with the problems with monitoring data that exhibit autoregression and the surveillance of multivariate data. Additional problems not discussed here concerns e.g. those associated with the data quality (reporting delays, measurement errors and revisions).

3.2.1 Effect of smoothing

In surveillance it is important to have control over the false alarms by being able to make a statement regarding how often we can expect a false alarm by a certain surveillance system. The variability of the process affects the false alarm rate. In order to reduce the false alarm rate some authors, e.g. Koskinen and Öller (2003), recommend that the observations should be smoothed after differentiation, see also Öller (1986). One motivation of the smoothing, has been the reduction of white noise and by that the false alarm rate. Another motivation has been a time adjustment when using multivariate data, where the turning points of the different processes are not always the same.

Smoothing by kernel estimators is used by e.g. Hall et al. (1995). Often the differentiated observations $y(t)$ are smoothed according to

$$\widetilde{Y}(t) = \lambda Y(t) + (1-\lambda)\widetilde{Y}(t-1),$$

where $0<\lambda<1$ can be determined in different ways. Koskinen and Öller (2003) estimate the transition probabilities and the smoothing parameter simultaneously from historical information, with the criterion of minimizing a cost-function based on the sum of two measures of error (the Brier probability score and the proportion of wrongly classified states) obtained by classifying observations from a previous period of data. The Brier probability score, also referred to as the Quadratic probability score, is the mean square error for the posterior probability, i.e. the average squared deviation between the true state (0 or 1) and the posterior probability.

The smoothing of observations reduces the variance and hence reduces the false alarm probability. However, there are also disadvantages in a surveillance situation as will be seen in the Monte Carlo study on the effect of smoothing on the *HMlin* method. Results will be given on the distribution of the alarms, both conditional of no change and of a turn, and the probability of detecting a turning point within a specified time.

Data on the (logarithm of the) Swedish IP was used to get a reasonable simulation model (see Appendix 1). A linear function was fitted to the officially dated (National Institute of Economic Research (1992)) expansion phase 1987Q2 to 1989Q3. The observations on X, under event D, are simulated using the following model

$$X^D(t) = \mu^D(t) + \varepsilon(t), \qquad (13)$$

where $\mu^D(t) = 11.194 + 0.0069 \cdot t$ and $\varepsilon(t) \sim$ iid N[0; 0.016]. In order to evaluate the alarm properties under event C the following model is used:

$$X^C(t) = \mu^{C\tau}(t) + \varepsilon(t), \qquad (14)$$

where $\mu^{C\tau}(t) = 11.194 + 0.0069 \cdot t - 2D_1 \cdot 0.0069 \cdot (t-\tau+1)$, $t = \{1, 2, ...\}$, and

$$D_1 = \begin{cases} 1, & t \geq \tau \\ 0, & \text{otherwise} \end{cases}$$

and $\varepsilon(t) \sim$ iid N[0; 0.016].

Replicates of sequences of three expansion phases and four recession phases are simulated.

The alarm limit for the smoothed data (λ=0.3) is adjusted to yield the same MRL^0 as for the unsmoothed data (λ=1).

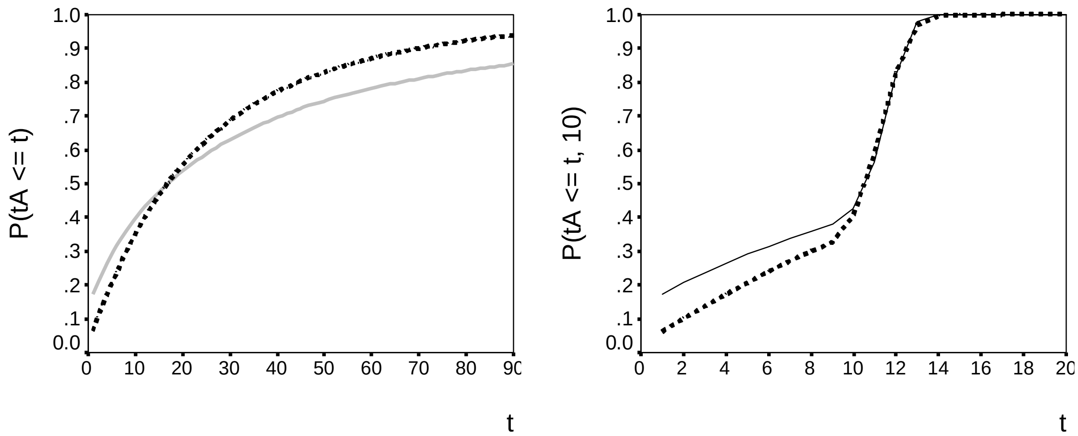

Figure 2. HMlin method, λ=1 (···) and λ=0.3 (——). Left: Distribution of the time of an alarm, conditional on event D (no turn). Right: Distribution of the time of an alarm, conditional on τ=10.

The common value of the MRL^0 is 17. For both cases, λ={1, 0.3} we have a large enough sample so that the standard error of the median is less than 0.15. We can see in Figure 2 that the density of the time of the alarm has different skewness for λ=1 and λ=0.3.

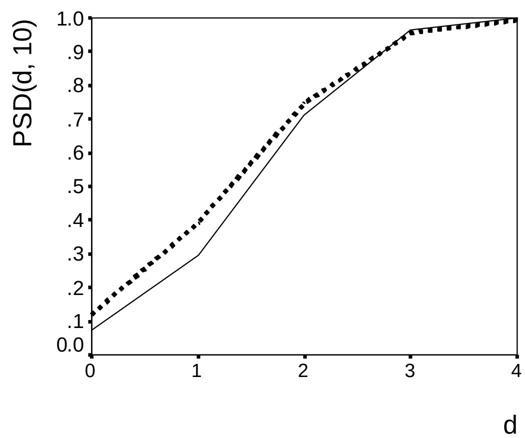

Figure 3. Probability of successful detection within d time points for τ=10. HMlin method, λ=1 (···) and λ=0.3 (——).

The probability of successful detection (PSD) is the probability of detecting a turn within d time units, that is

$$P\left((t_A - \tau) \le d \middle| t_A > \tau = \tau_0\right).$$

For both cases, λ={1, 0.3}, the number of replicates is large enough to yield a standard error of PSD of less than 0.0030. The reduced distinctness of the turning point, due to smoothing, decreases the probability of successful detection (see Figure 3).

3.2.2 Seasonal variation

The variables (leading economic indicators) are measured monthly or quarterly and thus often contain seasonal variation, which could complicate the monitoring. The seasonal variation can be considerable, as is seen in Figure 4.

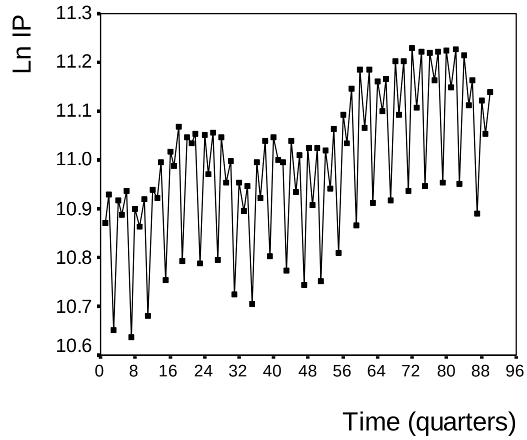

Figure 4. Swedish Industrial Production, quarterly data (1970Q1 to 1992Q2). Source: National Institute of Economic Research, Sweden.

If seasonality is neglected in the monitoring, it could lead to seriously wrong conclusions. In the monitoring situation here, it is important that the time of the turns is preserved after the adjustment in order to make actions induced by an alarm powerful. The problem of altered change points by seasonal adjustment has been briefly discussed in a fixed sample context, see e.g. Ghysel and Perron (1996) and Franses and Paap (1999).

Often the monitoring is made on data that have been adjusted for seasonality. The question of whether the seasonal variation can be considered stable over time is treated by e.g. Canova and Hansen (1995) and Busetti and Harvey (2003). These issues are important to consider when deciding the method for seasonal adjustment. The effect of using different filters in order to adjust for seasonality is analyzed in Andersson and Bock (2001) and it is demonstrated that the detection of a turn is delayed when data are differentiated or when a one-sided moving average is used. In the monitoring, the largest reduction in probability of detection is caused by the moving average.

3.2.3 Autoregression

Economic time series often exhibit strong autocorrelation. For instance, this can this be a problem when the sampling intervals are short (Luceno and Box (2000)). Lahiri and Wang (1994) evaluate the performance of a monitoring system where a model with autoregressive errors is assumed and where the posterior probability is used together with an alarm limit. The same alarm limit is used for models with autoregressive errors of different orders. They find that the introduction of autoregression in the errors leads to a smaller forecast error within phases but increases the risk of wrong inference concerning turning points. Three ways to deal with the effect of autoregressive errors is to i) use the correct likelihood ratio alarm statistic, see Pettersson (1998), ii) ignore the dependency but to adjust the alarm limit or iii) to monitor the residuals of an estimated model of the dependency. If the assumption of an independent process is used, when it is in fact dependent over time, the result is an increased false alarm probability. The consequence of autoregression in the process is examined in the general theory of surveillance where also remedies are suggested. For a review, see Pettersson (1998) and Frisén (2003). The situation with multivariate time-dependent data, modeled using a VAR model, is studied in Pan and Jarrett (2004).

Ivanova et al. (2000) argue that the effect of the autoregressive parameters will largely be captured by the probabilities of remaining in the current state (p_{11} or p_{22}). Many of the suggested methods of surveillance assume that the possible autoregression is not a severe problem, as is also assumed here.

3.2.4 Adjusting for trend

Many macroeconomic variables can be characterized as cyclical movements around a trend. In order to distinguish the movements and make the time series stationary around the cycle it is sometimes necessary to adjust for the trend. In model (1) no separation between the trend and the cycle is made.

Adjusting for trend implies a data transformation, which may result in a distortion of the characteristics of the original series, whereby the surveillance will not give reliable signals.

The effect of trend removal has often been studied in a non-surveillance context, but the results can still be used in this discussion. Gordon (1997) studies the effect of trend removal for predictive densities of the US GDP and warns against using other information from the data than that which is directly associated with the business cycle turning points. Canova (1998) discusses trend removal and evaluates the effect using several different approaches, among them first order differentiating. One conclusion from the study is that linear trend removal does result in turning point times which do not correspond to the official turning point times of the National Bureau of Economic Research (NBER), USA. In another paper Canova (1999) points out that previous research has shown that the trend may interact with the cyclical component and is therefore difficult to isolate. The general conclusion is that statements concerning the turning points are not independent of the statistical assumptions needed to extract trends.

In most HMM approaches, and also the one considered here (*HMlin*), differentiation is used and the surveillance is made on the differentiated process, whereas in *SRlin* and *SRnp* , surveillance is made on the undifferentiated process. The removing of the trend has less effect on the possibility to distinguish the turning points when analyzing short time series. Since the *SRlin* and *SRnp* methods are applied to a part of the time series that contains one turning point at most, no attempt to separate the trend from the cycle is made.

3.2.5 Multivariate problems

In the common movement approach, a business cycle is characterized as the cyclical movement of many economical activities. This is one example of how multivariate data is used. The common movement approach demonstrates that important information is contained in the relation between the turns of different indices. This information can be utilized, either by transforming the problem to a univariate one (by using a composite index of leading indicators) or by applying another method for surveillance of multivariate data. We discuss general approaches in the theory of multivariate surveillance and review some important contributions to the multivariate approaches to business cycles.

Wessman (1998) demonstrates that when the variables have the same change point (or known time-lag) the minimal sufficient alarm statistic is univariate. An example is that several leading indices with a known time lag concerning the turning point can be the base for one leading economic index. Reducing data to a univariate index is, in fact, the recommendation also for most of the earlier studies.

Stock and Watson (1991) and Stock and Watson (1993) model the common movements of coincident variables as arising from an unobservable common factor (the overall state of the economy). The key elements are the selection of variables and the estimation of the common factor.

Diebold and Rudebusch (1996) consider the common movements of coincident variables where the common factor is assumed to be governed by a two-state HMM. Kim and Nelson (1998) use the same approach as in Diebold and Rudebusch (1996) and find that the main cause of increase in forecast accuracy is the ability to capture the common movement among several variables instead of just one, whereas the prior assumptions concerning the transition probabilities has a minor influence.

Hamilton and Perez-Quiros (1996) compare the accuracy in predicting the phases of U.S. real gross national product using univariate and bivariate linear models, where the latter included a composite leading index (CLI), and corresponding HMM. Adding a CLI to the linear model was found to result in the greatest increase in accuracy, whereas using HMM makes no substantial increase in accuracy.

Koskinen and Öller (2003) utilize multivariate information by monitoring a joint vector of leading indicators with a common time of turn.

Birchenhall et al. (1999) exploit the feature of a business cycle, of common movements across variables, by extracting a business-cycle index from a vector of time series. As in the works by Stock and Watson, the selection of variables is an important element.

One alternative to the transformation to surveillance of a univariate index is to base the multivariate surveillance on the union intersection principle for the marginal processes. This is for example done in Woodall and Ncube (1985), where one surveillance system is applied to each process, and the aim is to detect the first change in any of the processes. Kalgonda and Kulkarni (2004) propose a surveillance method for a VAR process, where each process is standardized and the alarm statistic is the maximum of the standardized processes. Kontolemis (2001) compares turning point identification based on individual series to a multivariate approach (in all cases HMM approaches). It is shown that the business cycle chronology based on the latter approach is closer to that of NBER than the turning points obtained from individual series.

An optimal multivariate LR method can be based on the joint density function (Sonesson and Frisén (2004)). For reviews on the theory of multivariate surveillance, see e.g. Basseville and Nikiforov (1993), Lowry and Montgomery (1995), Ryan (2000) or Sonesson and Frisén (2004).

4. A comparison of some likelihood based approaches for detecting a turn in a leading index

When evaluating a surveillance method it is important to be able to make statements regarding the statistical properties of the method, for example the false alarm rate and the expected delay time for a motivated alarm. In Section 4.1 the performance of the methods are demonstrated for a turn in the Swedish industrial production and in Section 4.2 statistical properties of the three methods (*SRlin*, *SRnp*, *HMlin*) are summarized.

4.1 Evaluation by data on the Swedish Industrial Production

The most common way to evaluate methods for the detection of turning points in business cycles is by using one set of data. Here we use quarterly data on the Swedish industrial production. The three methods (*SRnp*, *HMlin* and *SRlin*) are applied to the period 1987Q2 to 1992Q2.) .

According to official records (National Institute of Economic Research (1992)), the period contains one turn, a peak, at time 1989Q3 ($t = 10$) which implies that the time of change is 1989Q4 (τ=11). The period is displayed in Figure 5. The official turning point times can often be based on more information than data. This other information might make the official time different than it should have been, if only the IP data had been used. Figure 5 indicates that the turning point in the data is earlier than the official time. Thus the methods, using only the IP data, can not be expected to be good at indicating the recorded official time for this realization. All three methods give alarms earlier than the official times for this set of data.

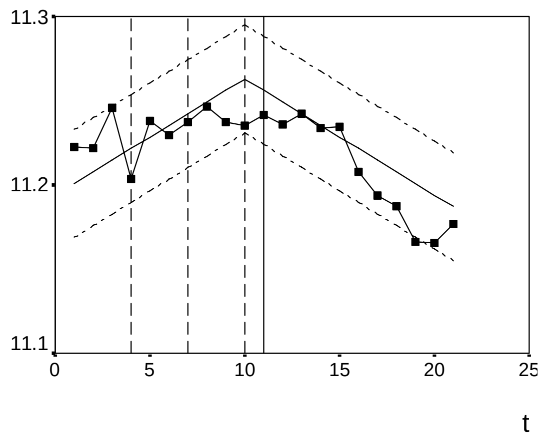

Figure 5. Seasonally adjusted values of the (logarithm of) Swedish industrial production, for the period 1987Q2:1992Q2. The official time of change, 11, is marked with a solid vertical line. The alarm times (t_A={4, 7, 10} for SRnp, HMlin and SRlin, respectively) are marked with dashed vertical lines. The model for μ with τ=11, used by SRlin and HMlin, is marked with a solid curve (see Section 3.2.1). The 2.5th and 97.5th percentiles of values of the observations according to the model is marked as dotted curves. Source: National Institute of Economic Research, Sweden.

Both the *SRlin* method and the *HMlin* method use the assumption of a piecewise linear model for μ (see (3)). The piecewise linear model fits less well at the turning point as we have a plateau. McQueen and Thorley (1993) argue that it is reasonable that recessions tend to be preceded by plateaus. A plateau will result in a tendency to give alarms just before the turn. It can be discussed whether this is a drawback or not. An early indication of a coming recession is a plateau. In this light, alarms just before the turning point can be considered to be good, since they can be seen as warnings.

The two parametric methods, *SRlin* and *HMlin*, rely on the assumption that the parameters (slopes and standard deviations) are known or possible to estimate with great certainty. Here we use the data on Swedish IP from the period 1970Q1:1987Q1 and the estimation procedure described in Appendix 1 to estimate the parameters. The resulting signal-noise ratios are: $\hat{\beta}_1 / \hat{\sigma}_1$ = 0.47 (expansion phase) and $\hat{\beta}_2 / \hat{\sigma}_2$ = 0.40 (recession phase). These estimates are used for the *HMlin* method when calculating the posterior probability. The *SRlin* method assumes a symmetric turning point and homoscedasticity.

We use pooled (by their frequency) estimates for both parameters, resulting in $\hat{\beta}/\hat{\sigma}=$ 0.41, when calculating the alarm statistic. The HMlin method also includes two transition probabilities, whose values are estimated according to (12).

An actual data set, as used above, is representative of the specific time period and situation at hand. However, the real data set might deviate stochastically from the process of interest. When a method is intended to be used on future data, then it is the properties of the process that are important for that method. In order to be able to make a general statement about, for example, the average delay time of a method, it is necessary to replicate the performance of the method at a turning point. Then Monte Carlo methods are a valuable tool.

4.2 Statistical properties

In Andersson et al. (2004) an evaluation was made of the properties of the three methods by means of a Monte Carlo study. The major results are summarized here. The evaluation was made both for the situation when the correct parameter values were used and for the situation when the values were misspecified.

4.2.1 Correct parameter values used

For all three methods the alarm limits are set to yield $\text{MRL}^0 = 17$ (quarters), but the distribution of the alarm times differs. The *HMlin* method has more frequent alarms at early time points, but low alarm probability later on, compared to the others.

The expected delay measures the delay time for a motivated alarm (i.e. the time between τ and t_A). This measure depends on τ. For a turning point within a year from the last turn, the *SRnp* method has the longest expected delay. For turning points that occur later, the *SRnp* method has a (slightly) shorter delay than the *HMlin* method. *SRlin* has the shortest delay time for every value of τ.

In a practical situation it is important to have a strategy for what action to take when a turning point is signalled. The predictive value (see Frisén (1992)) reflects the trust you should have in an alarm. For the *HMlin* method the high alarm probability at the first time point results in that alarms at t=1 are of little value, whereas the predictive value for *SRnp* and *SRlin* at this time have much higher predictive values. The alarms that come at time points t=4 and hence forward have predictive values of (at least) 0.75 for all three methods.

4.2.2 Incorrect parameter values used

The effect of using wrong parameter values for μ was evaluated both for the situation when only the slope after the turn was misspecified and for the situation when both slopes (pre-turn and post-turn) were misspecified. Also in this evaluation the alarm limits are set to yield $\text{MRL}^0 = 17$ (quarters).

When only the post-turn slope is misspecified, it has very little effect on the conditional expected delay and the predictive value.

When both slopes (pre-turn and post-turn) are misspecified, the effect on the conditional expected delay and the predictive value is major. For small and moderate values of τ, (τ<10), the delay time is longer when the slopes are specified as being too steep. Thereafter the delay time is shorter for misspecified slopes. The price for the short delay times is however that the predictive value of those alarms is low.

In view of these results, using a method that does not require any parametric values for μ is a safe way, particularly since the properties of the *SRnp* method are almost as good as those of the *SRlin* method.

5. Discussion

When estimating the parameters of the monitoring system, historical data is often used. Then there has to be a balance between, on one hand, the risk of using data sets that are too small, which results in estimates with a large variation and, on the other hand, the risk of using historical data which might be out of date. The user of a system for on-line detection is faced with the paradox that the parameters in the surveillance system might be estimated using previous data, which means that it is assumed that previous patterns will repeat themselves. However, the aim of the surveillance is to detect changes and by estimating parameters from previous data, the ability to detect changes in the current cycle might be diminished. Sarlan (2001) examines the change in intensity and duration of US business cycles and concludes that the modern business cycle is different from the historical one.

Other issues that must be considered in on-line detection is how to handle seasonal variation and whether to apply different kinds of transformations of data (e.g. trend removal and smoothing).

Since the surveillance methods are based on the ratio between the likelihoods given that the cycle has or has not reached a turn suitable models for the cyclical process must be found. Parametric models contain information, which should be used whenever it is reliable. However, wrong specifications do cause bad stochastic properties and misleading results. Here, a non-parametric approach, which works also when such reliable information is not available, is considered. The safe way with the non parametric method might be preferred in order to avoid risks of misleading results.

Many alarm systems, for example methods based on a Hidden Markov model, use assumptions regarding the intensity of the change (the transition probability). These parameters also need to be estimated. One approach is to use only the observed transition frequencies. Another approach is to estimate several parameters simultaneously. If, for example, both transition probabilities and smoothing parameters are estimated simultaneously then the parameters compensate for each other (a heavy smoothing is combined with a high transition probability). Thus it might be difficult to interpret the parameters separately. Different criteria can be used in the estimation and if the Brier probability score is used as a criterion when estimating the transition probabilities, it must be borne in mind that this measure does not take into account the order of the observations. As a result, the transition probabilities might again be difficult to interpret.

In many systems for on-line detection, it is assumed that the transition probability is constant, which implies that the time of the turn has a geometric distribution. This might not be in accordance with reality for business cycles, but can be interpreted as a way of avoiding to use strong assumptions regarding the intensity of turns. Good estimates of the transition probabilities are useful if the pattern is constant over time and will remain the same, even in the future. Technically, the inclusion of transition probability estimates in the monitoring system is easily done by likelihood ratio methods. However, it is important that the monitoring system has the ability to detect a turning point also when this happens at an unexpected time. Thus, it might be preferred to use a non-informative prior for the time of the turn in the suggested *SRlin* and *SRnp* methods, so as to avoid the risk of errors due to wrong assumptions or uncertain estimates.

The smoothing of the observations before applying a method of monitoring will reduce the variation and hence reduce the false alarm probability. However, the smoothing will introduce autocorrelation and, as pointed out by Öller and Tallbom (1996), will lead to a delayed signal. In this paper it is confirmed that the expected delay of a motivated alarm is increased by smoothing (when the false alarms are controlled). One method of surveillance, where the smoothing procedure is included in the alarm statistic and not made separately, is the *EWMA* method (see e.g. Crowder (1987), Domangue and Patch (1991) and Frisén and Sonesson (2004)). This method allows for a controlled false alarm rate at the same time as the variability is reduced by smoothing.

Economic time series often exhibit seasonal variation. Most data-driven filters can seriously alter the turning point times (see e.g. Andersson and Bock (2001)). Thus information from historical data or other prior knowledge, which makes the seasonal adjustment independent of the data to be monitored, is very valuable.

In surveillance of multivariate data, different approaches have been suggested. A common approach is to reduce data to a univariate index (e.g. a mean) and then apply methods for univariate surveillance to the index. Reducing data to a univariate index is, in fact, the recommendation also for most of the earlier studies. It can also be theoretically motivated if the turns occur at the same time (or with a known time lag). Surveillance with several leading indicators is, as it appears, an interesting topic for future research. Aspects influencing which approaches are optimal are the dependency structure of the variables (X_1, X_2, ...) as well as the dependency between the turning point times (τ_1, τ_2, ...).

When evaluating methods for on-line detection, it is important to consider the timeliness of the alarms. The rate of false alarms in surveillance is often controlled by setting the alarm limit so that the average time to the first false alarm is fixed (ARL^0 or MRL^0). A hybrid between surveillance and hypothesis testing is to set the alarm limit so that the probability of a false alarm is, at every time point, less than e.g. 0.05, which has been suggested by e.g. Chu et al. (1996). This approach has the disadvantage that the detection ability is very low if the change occurs after a long time (see Frisén (2003) and Bock (2004)).

The performance of a surveillance method should be measured by how quickly an actual change is detected (for example by the expected delay of an alarm or by the probability of successful detection). Evaluations using the power or using measures like the MSE does not take the timeliness into account.

Once the evaluation measures have been determined, there are still two main roads for evaluating a method of surveillance. In this paper, the effect of smoothing is evaluated by in a simulation study, whereas the monitoring methods are evaluated with a set of real data; a period of the Swedish IP. Evaluation of the properties of a method by one sample of real data is difficult. One difficulty is to know whether the turnout of the sample is typical. Evaluations by several real data sets (instead of just one) would decrease some of the stochastic variation in the measures of evaluation. However, if these analyses are not totally independent (for example if the same parameter estimates are used) then some of the stochastic components would keep their variance.

The results from the application of the three methods on a period of the Swedish IP did not contradict the conclusion by the statistical properties that the non-parametric method is a safe way without much loss of efficiency.

APPENDIX 1. Model for Monte Carlo study of the effect of smoothing

Data on the Swedish IP was used to get a reasonable simulation model. For the dating of recession and expansion phases, official records are used (National Institute of Economic Research (1992)),

The data on ln(IP) is seasonally adjusted and then, for each expansion and recession phase, a polynomial regression function is fitted.

$$y_{ij}(t) = \mu_{ij}(t) + \varepsilon_{ij}(t) = (\theta_{0ij} + \theta_{1ij} \cdot t + \theta_{2ij} \cdot t^2) + \varepsilon_{ij}(t) ,$$

where i = {expansion} or {recession} and j = {1, 2, 3, 4}.

Results for the intercept-adjusted polynomials are given in Table 1.

Table 1. Simulation model for the estimation period

i	j	θ_0	θ_1	θ_2	sd[$\varepsilon(t)$]
Expansion	1	10.707	0.023	-0.0002	0.004
	2	11.056	-0.019	0.0005	0.024
	3	12.678	-0.075	0.0008	0.013
Recession	1	10.920	0.009	-0.0009	0.007
	2	10.550	0.046	-0.001	0.018
	3	8.426	0.117	-0.001	0.013
	4	11.615	-0.016	0.0001	0.020

Acknowledgement

Constructive comments by Professor Lars-Erik Öller, Statistics Sweden, are gratefully acknowledged. This work was supported by the Swedish Council for Research in the Humanities and Social Sciences and by the Bank of Sweden Tercentenary Foundation.

References

Andersson, E. (2002) Monitoring cyclical processes - a nonparametric approach, **Journal of Applied Statistics**, 29, (7) pp. 973-990.

Andersson, E. (2004). Influence of turning point characteristics on a system for non-parametric surveillance. VIII International Workshop on Intelligent statistical quality control ISQC '04, Warsaw, Poland, Publishing House of Warsaw School of Information Technology.

Andersson, E. and Bock, D. (2001) On seasonal filters and monotonicity, **Research report** 2001:4, Department of Statistics, Göteborg University, Sweden.

Andersson, E., Bock, D. and Frisén, M. (2004) Statistical surveillance of cyclical processes with application to turns in business cycles, **Accepted for publication in Journal of Forecasting**pp.

Basseville, M. and Nikiforov, I. (1993) **Detection of Abrupt changes- Theory and Application** *(*Englewood Cliffs, Prentice Hall).

Birchenhall, C. R., Jessen, H., Osborn, D. R. and Simpson, P. (1999) Predicting US business-cycle regimes, **Journal of Business & Economic Statistics**, 17, (3) pp. 313-323.

Blondell, D., Hoang, H., Powell, J. G. and Shi, J. (2002) Detection of financial time series turning points: A new CUSUM aproach applied to IPO cycles, **Review of Quantitative Finance and Accounting**, 18, pp. 293-315.

Bock, D. (2004) Aspects on the control of false alarms in statistical surveillance and the impact on the return of financial decision systems, **Research report**, Statistical Research Unit, Göteborg University, Sweden.

Busetti, F. and Harvey, A. C. (2003) Seasonality tests, **Journal of Business and Economic Statistics**, 21, (3) pp. 420-436.

Canova, F. (1998) Detrending and business cycle facts, **Journal of Monetary Economics**, 41, pp. 475-512.

Canova, F. (1999) Does detrending matter for the determination of the reference cycle and the selection of turning points?, **The Economic Journal**, 109, (January) pp. 126-150.

Canova, F. and Ciccarelli, M. (2004) Forecasting and turning point predictions in a Bayesian panel VAR model, **Journal of Econometrics**, 120, (2) pp. 327-359.

Canova, F. and Hansen, B. E. (1995) Are seasonal patterns constant over time? A test for seasonal stability, **Journal of Business and Economic Statistics**, 13, (3) pp. 237-253.

Chu, C.-S. J., Stinchcombe, M. and White, H. (1996) Monitoring structural change, **Econometrica**, 64, pp. 1045-1065.

Crowder, S. (1987) A simple method for studying run-length distributions of exponantially weighted moving average charts, **Technometrics**, 29, pp. 401-407.

Delgado, M. A. and Hidalgo, J. (2000) Nonparametric inference on structural breaks, **Journal of Econometrics**, 96, pp. 113-144.

Diebold, F. X., Lee, J.-H. and Weinbach, G. C. (1994) Regime Switching with Time-varying Transition Probabilites. **In:** C. HARGREAVES (Ed.), **Nonstationary Time Series Analysis** pp. 283-302 (Oxford, Oxford University Press).

Diebold, F. X. and Rudebusch, G. D. (1989) Scoring the leading indicators, **Journal of Business**, 62, pp. 369-91.
Diebold, F. X. and Rudebusch, G. D. (1996) Measuring business cycles: A modern perspective, **Review of Economics and Statistics**, 78, (1) pp. 67-77.
Domangue, R. and Patch, S. C. (1991) Some omnibus exponentially weighted moving average statistical process monitoring schemes, **Technometrics**, 33, pp. 299-313.
Fang, Y. and Zhang, J. (1999) Performance of control charts for autoregressive conditional heteroscedastic processes, **Journal of Applied Statistics**, 26, (6) pp. 701-714.
Filardo, A. J. (1994) Business-Cycle Phases and Their Transitional Dynamics, **Journal of Business & Economic Statistics**, 12, (3) pp. 299-308.
Filardo, A. J. and Gordon, S. F. (1998) Business cycle durations, **Journal of Econometrics**, 85, (1) pp. 99-123.
Franses, H. P. and Paap, R. (1999) Does Seasonality Influence the Dating of Business Cycle Turning Points?, **Journal of Macroeconomics**, 21, (1) pp. 79-92.
French, M. W. and Sichel, D. E. (1993) Cyclical Patterns in the Variance of Economic-Activity, **Journal of Business and Economic Statistics**, 11, (1) pp. 113-119.
Frisén, M. (1986) Unimodal regression, **The Statistician**, 35, (4) pp. 479-485.
Frisén, M. (1992) Evaluations of Methods for Statistical Surveillance, **Statistics in Medicine**, 11, pp. 1489-1502.
Frisén, M. (1994) Statistical Surveillance of Business Cycles, **Research report** 1994:1 (Revised 2000), Department of Statistics, Göteborg University, Sweden.
Frisén, M. (2003) Statistical surveillance. Optimality and methods., **International Statistical Review**, 71, (2) pp. 403-434.
Frisén, M. and de Maré, J. (1991) Optimal Surveillance, **Biometrika**, 78, pp. 271-80.
Frisén, M. and Sonesson, C. (2004) Optimal surveillance by exponentially weighted moving average methods. Submitted.
Frisén, M. and Wessman, P. (1999) Evaluations of likelihood ratio methods for surveillance: Differences and robustness, **Communications in Statistics. Simulations and Computations**, 28, (3) pp. 597-622.
Gan, F. F. (1993) An optimal design of EWMA control charts based on median run-length, **Journal of Statistical Computation and Simulation**, 45, (3-4) pp. 169-184.
Ghysel, E. and Perron, P. (1996) The Effect of Linear Filters on Dynamic Time Series With Structural Change, **Journal of Econometrics**, 70, pp. 69-97.
Gordon, S. (1997) Stochastic trends, deterministic trends, and business cycle turning points, **Journal of Applied Econometrics**, 12, pp. 411-434.
Hackl, P. and Westlund, A. H. (1989) Statistical analysis of "Structural Change", **Empirical Economics**, 14, pp. 167-192.
Hall, P., Marron, J. S. and Titterington, D. M. (1995) On partial local smoothing rules for curve estimation, **Biometrika**, 82, (3) pp. 575-587.
Hamilton, J. D. (1989) A new approach to the economic analysis of nonstationary time series and the business cycle, **Econometrica**, 57, pp. 357-384.
Hamilton, J. D. and Perez-Quiros, G. (1996) What Do the Leading Indicators Lead?, **Journal of Business**, 69, (1) pp. 27-49.
Hawkins, D. L. (1992) Detecting shifts in functions of multivariate location and covariance parameters, **Journal of Statistical Planning and Inference**, 33, pp. 233-244.
Ivanova, D., Lahiri, K. and Seitz, F. (2000) Interest rate spreads as predictors of German inflation and business cycles, **International Journal of Forecasting**, 16, (1) pp. 39-58.

Jun, D. B. and Joo, Y. J. (1993) Predicting turning points in business cycles by detection of slope changes in the leading composite index, **Journal of Forecasting**, 12, pp. 197-213.

Kalgonda, A. A. and Kulkarni, S. R. (2004) Multivariate quality control chart for autocorrelated processes, **Journal of Applied Statistics**, 31, (3) pp. 317-327.

Kenett, R. and Pollak, M. (1996) Data-analytic aspects of the Shiryaev-Roberts control chart: surveillance of a non-homogeneous Poisson process, **Journal of Applied Statistics**, 23, pp. 125-137.

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, (2) pp. 188-201.

Kontolemis, Z. G. (2001) Analysis of the US business cycle with a vector-Markov-switching model, **Journal of Forecasting**, 20, (1) pp. 47-61.

Koskinen, L. and Öller, L.-E. (2003) A Classifying Procedure for Signalling Turning Points, **Journal of Forecasting**, 23, (3) pp. 197-214.

Lahiri, K. and Wang, J. G. (1994) Predicting cyclical turning points with a leading index in a Markov Switching model, **Journal of Forecasting**, 13, pp. 245-263.

Lai, T. L. (1995) Sequential Changepoint Detection in Quality-Control and Dynamical-Systems, **Journal of the Royal Statistical Society B**, 57, (4) pp. 613-658.

Layton, A. P. (1996) Dating and predicting phase changes in the U.S. business cycle, **International Journal of Forecasting**, 12, pp. 417-428.

Layton, A. P. (1998) A further test of the influence of leading indicators on the probability of US business cycle phase shifts, **International Journal of Forecasting**, 14, (1) pp. 63-17.

Layton, A. P. and Katsuura, M. (2001) Comparison of regime switching, probit and logit models in dating and forecasting US business cycles, **International Journal of Forecasting**, 17, (3) pp. 403-417.

LeSage, J. P. (1991) Analysis and development of leading indicator using a Bayesian turning-points approach., **Journal of Business and Economic Statistics**, 4, pp. 305-316.

Li, D. T. and Dorfman, J. H. (1996) Predicting turning points through the integration of multiple models, **Journal of Business and Economic Statistics**, 14, (4) pp. 421-428.

Lowry, C. A. and Montgomery, D. C. (1995) A Review of Multivariate Control Charts, **IIE Transactions**, 27, (6) pp. 800-810.

Luceno, A. and Box, G. E. P. (2000) Influence of the sampling interval, decision limit and autocorrelation on the average run length in Cusum charts, **Journal of Applied Statistics**, 27, (2) pp. 177-183.

Maravelakis, P. E., Panaretos, J. and Psarakis, S. (2004) EWMA chart and measurement error, **Journal of Applied Statistics**, 31, (4) pp. 445-455.

McQueen, G. and Thorley, S. (1993) Asymmetric business cycle turning points, **Journal of Monetary Economics**, 31, pp. 341-362.

Mudambi, R. (1997) Estimating turning points using polynomial regression, **Journal of Applied Statistics**, 24, (6) pp. 723-732.

National Institute of Economic Research (1992) Konjunkturläget Maj 1992, National Institute of Economic Research, Stockholm, Sweden.

Neftci, S. (1982) Optimal prediction of cyclical downturns, **Journal of Economic Dynamics and Control**, 4, pp. 225-41.

Page, E. S. (1954) Continuous inspection schemes, **Biometrika**, 41, pp. 100-114.

Pan, X. and Jarrett, J. (2004) Applying state space to SPC: monitoring multivariate time series, **Journal of Applied Statistics**, 31, (4) pp. 397-418.

Pettersson, M. (1998) Evaluation of some methods for statistical surveillance of an autoregressive process, **Research report** 1998:4, Department of Statistics, Göteborg University, Sweden.
Roberts, S. W. (1959) Control Chart Tests Based on Geometric Moving Averages, **Technometrics**, 1, pp. 239-250.
Roberts, S. W. (1966) A Comparison of some Control Chart Procedures, **Technometrics**, 8, pp. 411-430.
Robertson, T., Wright, F. T. and Dykstra, R. L. (1988) **Order Restricted Statistical Inference** *(*John Wiley & Sons Ltd).
Ryan, T. P. (2000) **Statistical methods for quality improvement** *(*New York, Wiley).
Sarlan, H. (2001) Cyclical aspects of business cycle turning points, **International Journal of Forecasting**, 17, (3) pp. 369-382.
Shewhart, W. A. (1931) **Economic Control of Quality of Manufactured Product** *(*London, MacMillan and Co.).
Shiryaev, A. N. (1963) On optimum methods in quickest detection problems, **Theory of Probability and its Applications**, 8, pp. 22-46.
Sonesson, C. (2003) Evaluations of Some Exponentially Weighted Moving Average Methods, **Journal of Applied Statistics**, 30, (10) pp. 1115-1133.
Sonesson, C. and Frisén, M. (2004) Multivariate surveillance. **In:** A. LAWSON and K. KLEINMAN (Ed.), **Spatial surveillance for public health.** pp. (New York, Wiley).
Srivastava, M. S. and Wu, Y. (1993) Comparison of EWMA, CUSUM and Shiryayev-Roberts Procedures for Detecting a Shift in the Mean, **Annals of Statistics**, 21, pp. 645-670.
Stock, J. H. and Watson, M. W. (1991) A Probability Model of the Coincident Economic Indicators. **In:** K. LAHIRI and G. H. MOORE (Ed.), **Leading Economic Indicators: New Approaches and Forecasting Records** pp. 63-89 (Cambridge, Cambridge University Press).
Stock, J. H. and Watson, M. W. (1993) A Procedure for Predicting Recessions with Leading Indicators: Econometric Issues and Recent Experience. **In:** J. H. STOCK and M. W. WATSON (Ed.), **Business Cycles, Indicators and Forecasting** 28, pp. 95-156 (Chicago, University of Chicago Press).
Theodossiou, P. T. (1993) Predicting Shifts in the Mean of a Multivariate Time Series Process: An Application in Predicting Business Failures, **Journal of the American Statistical Association**, 88, (422) pp. 441-449.
Wessman, P. (1998) Some Principles for surveillance adopted for multivariate processes with a common change point, **Communications in Statistics. Theory and Methods**, 27, pp. 1143-1161.
Westlund, A. H. and Zackrisson, U. (1986) On the prediction of structurally varying systems, **Technological Forecasting and Social Change**, 30, pp. 63-72.
Wetherhill, G. B. and Brown, D. W. (1991) **Statistical Process Control: Theory and Practice** *(*London, Chapman and Hall).
Woodall, W. H. and Ncube, M. M. (1985) Multivariate Cusum Quality Control Procedures, **Technometrics**, 27, pp. 285-292.
Zarnowitz, V. and Moore, G. H. (1982) Sequential signals of recessions and recovery, **Journal of Business**, 55, pp. 57-85.
Zellner, A., Hong, C. and Min, C.-K. (1991) Forecasting Turning Points in International Output Growth Rates Using Bayesian Exponentially Weighted Autoregression, Time-Varying Parameter, and Pooling Techniques, **Journal of Econometrics**, 49, (1-2) pp. 275-304.
Zuehlke, T. W. (2003) Business cycle duration dependence reconsidered, **Journal of Business and Economic Statistics**, 21, (4) pp. 564-569.

Öller, L.-E. (1986) A note on exponentially smoothed seasonal differences, **Journal of Business and Economic Statistics**, 4, pp. 485-489.
Öller, L.-E. and Tallbom, C. (1996) Smooth and timely business cycle indictors for noisy Swedish data, **International Journal of Forecasting**, 12, pp. 389-402.