Sequential Detection of US Business Cycle Turning Points: Performances of Shiryayev-Roberts, CUSUM and EWMA Procedures

Bakhodir A Ergashev
Department of Economics
Washington University
St. Louis, MO, USA
e-mail: ergashev@economics.wustl.edu

First Draft: February 1, 2004
This Version: March 5, 2004

Abstract
In this paper we consider the problem of sequential detecting change points in economic time series. We compare the performances of three well known procedures, Shiryayev-Roberts, CUSUM and EWMA, in the problem of early detection of the US business cycle turning points using leading indicators or some financial series. The comparison was done separately for detecting recessions and expansions during the period of 1955-2003. We found that in most cases the Shiryayev-Roberts procedure is superior to the other two in detecting turning points with leading indicators. At the same time the CUSUM procedure performs better in detecting turning points with stock price indices.

1. Introduction

1.1. Instability in time series

Recent research shows that instability in time series is not uncommon. Stock and Watson (1996) assessed the prevalence of parameter instability in 76 monthly US economic time series. They rejected stability in most of the models built using the above-mentioned series. Therefore being able to find changes in time series helps one build more accurate models and make more adequate inferences. During the last two decades there has been significant progress in theoretical studies of the change point problems as well as in their applications. Without attempting to review the extensive literature dealing with the whole set of problems concerning change points we would like to emphasize several approaches that have been used in the recent literature dealing with change points in economic time series.

The first approach is *hypothesis testing*. Here H_0 states there is no change in the parameters of the underlying model during a given time period while the alternative hypotheses may differ from a simple one that states one change (in some known or unknown moment of time) to multiple changes in several parameters. Chow (1960) is the most well known test for structural change. However it has an important limitation in that the change point must be known a priori. Addressing this limitation, Quandt (1960) developed the likelihood ratio test that considers the time point with the largest Chow statistic as a possible change point. Brown, Durbin and Evan (1975) test for structural breaks in regression models using Page's (1954) CUSUM test[1] on the recursive residuals.

[1] See the next Sections for the definitions of the CUSUM statistic and recursive residuals.

The second approach deals with the problem of *estimation* of the unknown change point. Quandt's (1960) test is one of the most popular tests in change point estimation.
Bai (1994, 1997) develops a theory of least square estimation of a change point in multiple regression models. As examples of practical applications we mention Hansen (2001), who using Quandt's test found a break in US Labor productivity sometime between 1992 and 1996, and Zivot and Andrews (1992) who tested unit root against the alternative of stationarity with a structural change at some unknown time for the US GNP.

Recently Chib (1998) has proposed a new approach for the multiple change point modeling in the Bayesian framework. The changes occur at times when a state variable evolving according to a hidden Markov process jumps. Among the other features, the approach can fit models in which the change-point probability is not a constant and depends on the regime. It should be noted that the recent developments in the Monte Carlo Markov Chain (MCMC) models are making the Bayesian approach more attractive.

There is also a brand of the turning point literature that considers parameters themselves as stochastic processes. In this case every observation time becomes a change point (see, for example, Cooley and Prescott (1973, 1976)).

Note that the approaches mentioned above are retrospective. The historical data is used to find the possible change points in the past. Therefore in these settings one cannot approach the problem of predicting future change points. The main reason is that they do not take into account the sequential nature of the time series data. For example, we need a different setting to address the question of whether Leading Economic Indicators or some financial series can predict cyclical turning points. It is well known that simple rules, such as three consecutive-declines, do not help much. To answer questions like this, one needs more sophisticated techniques that are optimal in some sense. These types of problems are called the problems of e*arly detection/prediction of change points.* Here one uses the historical data to predict future changes in the time series.

In the next section we formulate the precise definition of *optimality* in the early detection problem. But the intuition behind optimality is simple: one wants to detect a change as early as possible, but does not want to have too often false alarms. There must be some tradeoff between these two conflicting desires. The concept of optimality here assumes that changes are not frequent and they occur suddenly at unknown times in the future.

The approach is general and it can be applied to different problems. In statistical quality control, this approach is used to detect changes in the quality of the manufacturing product as early as possible in order to minimize the losses connected with producing defective products. In economics, for example, it is of significant importance to detect turning points in business cycles as early as possible. Neftci (1982) was the first to introduce this approach to economics. In finance, for example, practitioners need to know when it is time to reconsider the coefficient estimations of financial models (such as CAPM, Fama and French's (1993) three factor model etc.) so that they have frequently updated (and therefore more accurate) models. But very frequent updates are costly. Thus, there is a need for the special rules that monitor the models by signaling times to update in the models' parameters (see Knoth, 2002).

1.2. Outline of paper

In this paper we consider the last problem of early detection/prediction of change points. We define change points as the moments in time when the parameters of the underlying model change. We assume that change events are not very frequent and the parameters are constant between two consecutive change points.

In Section 2 we formulate the problem and introduce three well known methods that can be used in detecting change points sequentially. Then in Section 3 we apply the three procedures to the early detection of the US business cycle turning points problem using the OECD's Composite Leading Indicators (CLI) Index as an observing series. One problem with using the OECD data is that it has recently been revised. Therefore, in order to see how our methodology performs when it is used for unrevised data, we applied it to the S&P Composite Index. Theses results are reported in Section 4.

Our approach is close to that of Neftci (1982). Neftci uses Shiryayev's (1963) optimal Bayesian detection rule, which requires the knowledge of a priori probability of a change happening. Here we use three different rules, one of them being the frequentist version of Shiryayev's rule (called the Shiryayev-Roberts procedure), which is available from the sequential statistical analysis literature. We also consider a broader time period that includes several turning points, and using the above-mentioned procedures we try to detect not only the recessions but the expansions as well.

As we will see later, the use of an index of leading indicators or stock price indices for early detection of change points raises several issues that challenge the assumptions of the problem setting regarding an observed series: homosckedasticity, normality, independence, and the assumption that the magnitude of a change is known. We address some of these challenges and outline the future research that has to deal with them. Then we conclude. Some details of the computations and the graphs are presented in the Appendix.

2. Change Point Detection Problem

2.1. Problem

Suppose that we observe sequentially the following sequence of random variables

$$X_1, X_2, ..., X_n, ...$$

Assume ν – is an unknown change point time so that $X_1, ..., X_{\nu-1}$ are i.i.d. N(0,1) and $X_\nu, X_{\nu+1}, X_{\nu+2}, ...$ are i.i.d. N(m,1) where m – is some known positive number. Note that $\nu = \infty$ means there is no change in distribution – all random variables are i.i.d. N(0,1).
Also note that here we assume without loss of generality that observed random variables have unit standard deviation.

We need a detection rule/procedure with the following two *properties:*

- it should detect a change as soon as possible after the change has occurred;
- it should not give frequent false alarms that a change has occurred (if the change has not occurred yet).

Since these two properties conflict with each other, the desire to detect a change as soon as possible leads to frequent alarms before a change (or between two consecutive changes). A possible compromise is to fix the average time between two false alarms (average run length, ARL) when there are not any changes and among those rules try to find the rule that minimizes average delay (AD) in detecting true change.

Consider any rule and denote τ the time when the rule announces a change (stopping or alarm time). The compromise discussed above could be formulated more formally as follows. Find a rule with a stopping time τ such that

$$ARL \equiv E(\tau \mid \nu = \infty) \geq T \,, \tag{1}$$

where T is a fixed number (T=50,100, etc.)

and $$AD \equiv \max_{k=1,2,\ldots} E(\tau - k \mid \tau \geq k) \quad \rightarrow \quad \min \tag{2}$$

Here *E(X|A)* means the conditional expectation of random variable *X* given event *A*.

There are three well known procedures that are frequently used in the above formulated problem. Now we introduce them and discuss their properties.

2.2. CUSUM procedure

The first procedure, cumulative sum (CUSUM), was proposed by Page (1954). This procedure is based on the following stopping rule

$\tau_P \equiv \min\{n: \max_{1\leq k\leq n}(S_n - S_k) \geq A\}$, where $S_k = \sum_{i=1}^{k} \log\{f_m(X_i)/f_0(X_i)\}$, $f_d(x)$ is the probability density function of distribution *N(d,1)* and $A > 0$ is some given number, call the threshold.

In other words, the rule stops the observations and announces that a change has occurred at first n such that $\max_{1\leq k\leq n}(S_n - S_k) \geq A$.

Taking into account the normality assumption we can rewrite the rule as follows

$$\tau_P \equiv \min\left\{n: \max_{1\leq k\leq n}\left[m\left(\sum_{i=1}^{n} X_i - \sum_{i=1}^{k} X_i\right) - (n-k)\frac{m^2}{2}\right] \geq A\right\}.$$

Lorden (1971) proved that the CUSUM procedure is asymptotically optimal in the first order.

Since we do not know accurate approximations to ARLs of the CUSUM procedure we performed the Monte Carlo simulations in order to find the appropriate levels of threshold A that lead to desired levels of ARLs. These results are reported in the next section.

2.3. Shiryayev-Roberts procedure

Shiryayev (1963) (see also Shiryayev (1978)) solved the problem, formulated in (1) and (2), in a Bayesian framework in continuous time. As a limit of Bayesian solutions he proposed the following rule

$\tau_S = \inf\{t: \; SH_t(m) \geq B\}$, where $SH_t(m) = \int_0^t \exp\{W(t) - W(s) - m^2(t-s)/2\}ds$ and $W(t)$ is a standard Wiener process. Assuming stationarity of SH_t Shiryayev was able to derive the explicit expressions for its ARL and AD.

The following discrete analogue of this rule was considered by Roberts (1966):

$$\tau_R = \min\{n: \ R_n(m) \geq B\}, \text{ where } R_n(m) = \sum_{k=1}^{n} \prod_{i=k}^{n} [f_m(X_i)/f_0(X_i)].$$

Hereafter we will drop m and some other indices in order to present the formulas in a more readable format.

Due to the recursive nature of the R statistic it can also be calculated recursively as follows

$$R_n = (R_{n-1} + 1) \cdot \frac{f_m(X_n)}{f_0(X_n)}, \qquad n = 1,2,..., \quad R_0 = 0.$$

Under the assumption of normality this recursive formula can be rewritten as

$$R_n = (R_{n-1} + 1) \cdot \exp\{ mX_n - \frac{m^2}{2}\}.$$

Pollak (1985) proved that with randomization in the beginning of observations this rule is optimal. Following the convention we call this rule the Shiryayev-Roberts procedure.

Under the assumption that there is no change in their distributions ($\nu = \infty$) the process $SH_t - t$ and the sequence $R_n - n$ are martingales (see, for example, Pollak and Siegmund (1985, 1991)).

Since a Wiener process has a continuous trajectory it is obvious that $SH_\tau = B$. Therefore in continuous time (due to the optional stopping theorem)

$$ARL \equiv E(\tau_S \mid \nu = \infty) = E(SH_\tau \mid \nu = \infty) = B.$$

This formula simplifies ARL computations in continuous time. But in practical applications one has to deal with discrete time. Unfortunately in this case $ARL \neq B.$ This happens because of jumps of the sequence R_n at discrete times: we have $R_\tau = B + \chi$, where χ is the excess of R_n over B at time $n = \tau_R$. Here we use the following approximation proposed by Pollak and Siegmund (1991) in their formula (2) (see also Chapter 10 of Siegmund (1985)):

$$ARL \equiv E(\tau_R \mid \nu = \infty) \cong B / \exp(-0.538m). \qquad (3)$$

The Monte Carlo simulations reported in the next section show that this formula is quite accurate.

2.4. Exponentially Weighted Moving Average (EWMA)

The third procedure, the Exponentially Weighted Moving Average (EWMA), is very popular in quality control applications. It also has been applied in economics and finance. It has been used and recommended as one of the reliable rules for forecasting the volatility of asset prices by JP Morgan's (1996) RiskMetrics Group. Financial participants also use the EWMA for Value at Risk (VaR) estimations of stock markets (see, for example, Guermat and Harris (2001)). There are cases discussed in the literature where the EWMA was used for monitoring a Capital Asset Pricing Model (CAPM) (Knoth 2002) and forecasting exchange rates (see, for example, Muller-Plantenberg (2003)).

Perhaps one of the reasons for the popularity of EWMA is the simplicity of the EWMA statistic. At the same time, it should be noted that the EWMA procedure is defined by two parameters (λ and L defined below). Therefore for any given level of ARL there will be many choices of these two parameters with the same ARL but different ADs. Thus, an additional problem of *optimal design of EWMA* arises: for any given level of ARL one has to choose these two parameters so that AD is minimized for that choice. More formally, for any given level T of ARL optimal design means choosing λ and L so that

$$ARL(\lambda, L) = T \quad \text{and} \quad AD(\lambda, L) \to \min.$$

From this point we will work with optimally designed EWMA procedures and by EWMA we always mean optimally designed EWMA.
The stopping time of the EWMA rule is defined as

$$\tau_{MA} = \min\{n: \ |Y_n| \geq A\}, \text{ where } Y_n = (1-\lambda)Y_{n-1} + \lambda X_n, \ n = 1,2,..., \qquad Y_0 = 0, \tag{4}$$

and $0 < \lambda < 1$ is some constant called the smoothing parameter.
The control limit A is usually expressed in terms of the asymptotic standard deviation $\sigma = \sigma(\lambda)$ of the control statistic Y_t. That is $A = L\sigma$ for some $L > 0$, and the process is considered out of control whenever $|Y_t| > L\sigma$. It is easy to see that $\sigma^2 = \sqrt{\lambda/(2-\lambda)}$.
Due to the symmetric nature of the EWMA stopping rule τ_{MA}, one can use it to detect shifts in the mean in both directions.

In order to design an EWMA optimally one has to be able to compute its ARLs and ADs. Crowder (1987) presented the integral equations approach to this problem. Later Lucas and Saccucci (1990) used the Markov property of the EWMA statistic and they were able to compute its ARLs and ADs more accurately. Their approach also allowed them to design the EWMA optimally. At the same time it is well known that the algorithm they have proposed becomes less efficient and less accurate for small values of λ that are useful in detecting small shifts.

In Ergashev (2003) we presented some approximations to EWMA characteristics (ARL and AD) based on Novikov's (1990) exact formulas for the ARL and AD of the EWMA procedure. See Ergashev (2003) for the details. The accuracy of the ARL approximations presented there is verified to be satisfactory for ARLs not less than 100. But in this paper we would like to use ARLs as low as 50. Although in this case the ARL approximations are still good, the AD approximations become unsatisfactory. Therefore here we use Monte Carlo experiments in order to estimate the ADs of the EWMA procedures with small ARLs.

2.5. Some theoretical and numerical comparisons of the three procedures
In practical applications of the above defined procedures one has to take into account the fact that the value of the true shift in the mean (denote it μ) can be different from the anticipated shift, m. This fact plays an important role in practical comparisons. As we have mentioned earlier, the Shiryayev-Roberts and the CUSUM procedures are asymptotically optimal in sequentially detecting change points when μ=m. Pollak and Siegmund (1985) performed some numerical comparisons of these two procedures. They found that they are almost indistinguishable when the true shift is equal to the anticipated one (μ=m). But for larger values of μ the CUSUM does slightly better while for its smaller values the Shiryayev-Roberts rule seems preferable.

Although the EWMA is not optimal in the sense defined in (1) and (2), it is well known from quality control applications that it does not lose much (Novikov (1990), Srivastava and Wu (1993)). Srivastava and Wu (1993) compared the performances of all three rules and found that the EWMA rule performs better than CUSUM for small ARLs (around 100) in detecting small shifts. Their comparisons show that for small ARL values the Shiryayev-Roberts procedure becomes much better than the CUSUM procedure when μ<1.18m, while the CUSUM procedure is slightly better in the opposite case.

A numerical comparison of the EWMA and CUSUM rules by Lucas and Saccucci (1990) also showed that the EWMA rule is quite competitive in most practical situations (see their Table 6).

3. Can Leading Indicators predict Business Cycle Turning Points?

In his seminal paper, Neftci (1982) shows that the Index of Leading Indicators has a significant predictive power in predicting business cycle downturns and Shiryayev's Bayesian rule is able to detect them. In this section we apply the three other rules, Shiryayev-Roberts, CUSUM and EWMA procedures, to the same problem of early detection of US business cycle turning points by applying them to data that covers a broader time period.

3.1. Data and US business cycles

We use the OECD's Index of Composite Leading Indicators (CLI) available at http://www.oecd.org. This dataset covers the period between January 1955 and November 2003 resulting in 587 monthly observations. We use the above introduced three procedures in order to detect the change points in this series. If the leading indicators have any power to predict the business cycle turning points then the procedures should be able to detect them earlier. The following table describes the postwar US business cycle turning points – peaks and troughs.

Table 1. Postwar US business cycles

Peak	Trough	Duration, months	
		Between Peaks	Between Troughs
February 1945	October 1945		
November 1948	October 1949	45	48
July 1953	May 1954	56	55
August 1957	April 1958	49	47
April 1960	February 1961	31	34
December 1969	November 1970	117	118
November 1973	March 1975	47	52
January 1980	July 1980	74	64
July 1981	November 1982	17	27
July 1990	March 1991	109	100
March 2001	November 2001	129	129
Coverage	**Period**	**Average Duration**	
Our Sample	1955-2001	72	71
Past WWII	1945-2001	67	67

Source: NBER

3.2. An appropriate model for CLI

We could not reject the null hypothesis that the logarithms of CLI contain a unit root with the Augmented Dickey-Fuller test. The results are given in the next table.

Table 2. Augmented Dickey-Fuller unit root test[a]

$z_t = \mu + \gamma t + \rho_1 z_{t-1} + \rho_2 (z_{t-1} + z_{t-2}) + ... + \rho_k (z_{t-k+1} + z_{t-k}) + u_t$

Variable	Value	t-statistic
Sample Size	572	
k	14	
μ	0.019	2.11
γ	0.00001	1.87
ρ_1-1	-0.005	-1.98
ADF Statistic	-1.98	
5% critical value	-3.42	

[a] z_t represents the natural logarithms of the CLI data.

Since the residuals of the ADF test model had autocorrelations in some lags we performed a Phillips-Perron unit root test too. This test also was unable to reject unit root with the PP test statistic -0.725 while the 5% critical value is -2.867.

Since we could not reject the unit root hypothesis we model the logarithm of CLI, z_t, as

$z_t = c + z_{t-1} + \varepsilon(t)$, where $\varepsilon(t)$, $t = 1,2,...$ are the OLS residuals. (5)

3.3. Recursive residuals

Since it is natural to look at residuals to investigate departures from model specifications, in the model described above one can use its OLS residuals in order to detect turning points in CLI. But OLS residuals are not a very sensitive indicator of small or gradual changes in the parameters[2]. Also they are not independent and always sum to zero. Brown, Durbin and Evans (1975) advocated the use of recursive residuals. In their model the regressors were assumed to be non-stochastic. Under the null hypothesis that there are no changes in the model parameters they were able to prove that the recursive residuals are i.i.d. The model we consider here is autoregressive. Thus their result does not cover this case. We are not aware of studies that investigate the properties of recursive residuals in autoregressive regressions. Nevertheless we hope that recursive residuals are still sensitive and able to capture parameter changes in our model. The results reported in the next section show that this is the case.

Recursive residuals are defined as the standardized residuals from the regression of each observation at time t, the regression coefficients being calculated from the observations up to time t-1.

In our case, since the model is given by (4), we compute recursive residuals X_t as

$$X_t = [(z_t - z_{t-1}) - c(t-1)] / \hat{\sigma}(t-1), \quad (6)$$

where *c(t-1)* is the regression coefficient estimate based on the observations up to time t-1 inclusively and $\hat{\sigma}(t-1)$ is the standard deviation of $X_1, X_2, ..., X_{t-1}$.

[2] It should be noted that Ploberger and Kramer (1992) used OLS residuals in CUSUM test for a regression with non autoregressive independent variables.

3.4. Mean shifts in recursive residuals as turning points

If there are no changes in the parameters of the model (5) then the recursive residuals defined in (6) should have zero mean. We would like to detect changes in the parameters that lead to a shift in the mean of the recursive residuals. Mean shifts in positive direction correspond to the beginning of expansions while negative shifts correspond to the beginning of recessions. We use m=±1 as an appropriate shift in the mean of the recursive residuals we would like to detect. Since recursive residuals are standardized, this corresponds to one sigma shift in the mean in each direction. Later we discuss some other choices of the value of anticipated shift, m.

3.5. ARL computations

To be able to compare all three procedures we have to use the same ARLs for all three.
Following Neftci (1982) we use the average duration of the business cycles in our sample period as an appropriate level of ARL which is ARL=70 as is seen in Table 1. But we also consider some other choices given in the next two tables. In these tables we present the values of the thresholds of each procedure that correspond to a particular choice of ARL when m=1. For the CUSUM procedure we found them through a Monte Carlo experiment. The accuracy of the choices is verified in the third column of Table 3 by the Monte Carlo simulations. For the Shiryayev-Roberts procedure we used formula (3) in order to find the corresponding values of the threshold B. Again, the accuracy of the choices has been verified and it is reported in the third column of Table 3. The estimates of the AD for these two procedures are reported in the fourth column. They confirm the findings of Pollak and Siegmund (1985) and Srivastava and Wu (1993) that the Shiryayev-Roberts procedure is better than the CUSUM for small values of ARL when μ=m.

Table 3. Parameters of CUSUM and Shiryayev-Roberts procedures

CUSUM procedure			
Desired ARL	A (threshold)	ARL by Monte Carlo*	AD by Monte Carlo*
50	2.2	50.4	6.8
70	2.5	70.2	7.4
100	2.82	99.3	8.0
200	3.5	200.8	9.4
Shiryayev-Roberts procedure			
Desired ARL	B (found using (3))	ARL by Monte Carlo*	AD by Monte Carlo*
50	27.91	51.6	6.4
70	39.08	71.5	7.0
100	55.82	101.5	7.7
200	111.64	201.1	9.0

* - for each case 100,000 simulations have been performed.

In Table 4 we present the parameters of the EWMA procedures optimally designed to detect one sigma shift (m=1) in the mean value of the recursive residuals. The values of L and λ for ARL=100 are taken from Lucas and Saccucci's (1990) Table 4. For the other choices of ARL we found optimal L and λ the following way. First, we used the ARL approximation (6) from Ergashev (2003) in order to find different choices of L and λ with the same ARL. Then we used Monte Carlo simulations in order to find the choice with minimal AD. The results are given in Table 4.

Table 4. Optimal EWMA parameters in detecting one standard deviation shift

Desired ARL	L	λ	ARL by Monte Carlo*	AD by Monte Carlo*
50	2.100	0.23	51.3	6.7
70	2.209	0.20	70.9	7.3
100	2.298	0.16	100.7	8.0
200	2.584	0.16	201.5	9.3

* - for each case 100,000 simulations have been performed.

3.6. The results

In the next two tables we present our findings. Some details concerning the computations are given in the Appendix. Since all the three procedures performed similarly in all cases here we report the results for the case of ARL=70 only.

The Shiryayev-Roberts procedure was the best among the three in detecting recessions. It detected three recessions earlier than the CUSUM procedure. The EWMA procedure performed the worst. It failed to detect the beginning of the recession following April, 1960. See also graphs 1-3 of the Appendix.

Table 5. Detecting recessions with the three procedures

Peaks	**CUSUM**	**Shiryayev-Roberts**	**EWMA that resets at UCL***
August-57	January-56	January-56	January-56
April-60	August-59	August-59	Failed to detect
December-69	June-69	June-69	June-69
November-73	June-73	May-73	July-73
January-80	October-79	March-79	November-79
July-81	September-81	September-81	October-81
July-90	August-90	July-90	August-90
March-01	December-00	November-00	December-00

* - Since recessions are detected by shifts in the mean of the recursive residuals in negative direction (downward turns) this EWMA was designed to reset EWMA statistic Y_t to zero whenever it crosses the upper control limit (UCL). See the Appendix for more details.

In detecting the expansions all three procedures were a little later than in detecting the recessions. This happened probably because of the well known business cycle asymmetry: recessions are sharp and fast while expansions are mild and slow (Neftci (1984)).

In detecting expansions the Shiryayev-Roberts and the CUSUM procedures performed similarly while the EWMA's performance was not satisfactory; it failed to detect the last two expansions. See also graphs 4-6 in the Appendix. It should be noted that all three procedures are predicting a new expansion that began in September of 2003 according to the Shiryayev-Roberts procedure.

Table 6. Detecting expansions with the three procedures

Troughs	CUSUM	Shiryayev-Roberts	EWMA that resets at LCL*
April-58	May-58	May-58	May-58
February-61	April-61	April-61	May-61
November-70	January-71	January-71	January-71
March-75	May-75	May-75	June-75
July-80	July-80	July-80	July-80
November-82	February-83	February-83	February-83
March-91	June-91	June-91	Failed to detect
November-01	February-02	February-02	Failed to detect
New Expansion	October-03	September-03	November-03

* - Since expansions are detected by shifts in the mean of the recursive residuals in positive direction (upward turns), this EWMA was designed to reset EWMA statistic Y_t to zero whenever it crosses the lower control limit (LCL). See the Appendix for more details.

4. Do stock price indices have any predictive power in detecting US business cycle turning points?

Estrella and Mishkin (1998) found that stock price indexes have some predictive power in predicting the US recessions. If the methodology described in Sections 2 and 3 works it should reveal that predictive power.

Here we apply the methodology to the S&P Composite Index. The monthly data used here can be found on Robert Shiller's webpage http://www.econ.yale.edu/~shiller/. In order to have the comparable results we use the same procedure parameters as in the previous application (these parameter choices are described Section 3; see also the Appendix) and consider the same time period: January 1955-November 2003. Also, we use the same model (5) for the logarithm of the S&P Composite Index, z_t. This is a well know way of modeling stock prices.

Since the Shiryayev-Roberts and the EWMA procedures did not perform well we are not reporting their performances. The CUSUM procedure performed better than the other two, perhaps because of the prevalence of sudden and large shifts in the S&P series. This was not surprising since it is well known that the CUSUM procedure detects this kind of shifts better than the other two (see Section 2.5).

The results, presented in Graphs 7-8 in the Appendix, show that the stock price indexes have a moderate power in predicting US business cycles turning points.

While detecting the peaks - the beginnings of the recessions – the CUSUM procedure was able to detect on time the March 2001 peak only. It was late in detecting the November 1973 and July 1990 peaks, detecting them in December 1973 and October 1990, respectively. The other five peaks were not detected at all (see Graph 7).

The CUSUM procedure predicted the beginning of only two expansions out of 8 in our sample. It detected the March 1975 and November 1982 troughs in February 1975 and October 1982, respectively. And it was several months late in detecting the July 1980 trough (see Graph 8). But it should be noted that a choice of a smaller control limit could lead to slightly better results while increasing the rate of false alarms.

5. Discussion and Conclusions

In this paper we considered the problem of early detection of change points. We compared the performances of three well known procedures in detecting the US business cycle turning points. We considered a broader time period that included several turning points, and used the above mentioned procedures to detect not only recessions but expansions as well.
The results show that the Shiryayev-Roberts procedure performs better than the CUSUM and the EWMA procedures in detecting recessions using leading indicators. However, in detecting expansions the first two procedures performed similarly. The CUSUM procedure was the best among the three in detecting business cycle turning points using stock price indices. However, the results reported in Section 5 show that the predictive power of stock price indices is weak in predicting recessions, and is slightly better in predicting expansions.

The use of an index of leading indicators or some financial series in this regard raised several issues that challenge the assumptions of the problem setting regarding the observed series. The recursive residuals of the considered models were heteroskedastic, skewed and serially correlated. Also, in detecting the most turning points the so called misspecification problem arose; the true magnitude of a change in the mean of the recursive residuals was significantly different from the anticipated change of one standard deviation.

In order to better understand the properties of the two optimal procedures we applied them to the index of industrial production series, monthly records of which are available from 1930 at www.economagic.com. This study also confirmed that the Shiryayev-Roberts and the CUSUM procedures perform well even when some of the basic assumptions regarding the observed series are violated. Moreover, it confirmed that the CUSUM procedure is faster in detecting large and fast changes in the mean while the Shiryayev-Roberts procedure is faster in detecting small and gradual changes.

Further research needs to focus on dealing with the challenges to the problem assumptions mentioned earlier. For example, as a possible solution to the misspecification problem one could consider the problem of detecting a change point with multiple alternatives. In this study we assume that one knows the true value of a shift in mean, m. What if m is unknown and takes on of possible s (s=2,3,…) values? One solution would be considering the weighted sum of s Shiryayev-Roberts statistics that are designed to detect a particular shift. The asymptotic optimality of this procedure has been showed in Ergashev (1991). But in our sample this procedure was not able to outperform simple Shiryayev-Roberts procedure. Probably the Bayesian procedures with continuously updated weights are needed.

In conclusion, our study shows that the Shiryayev-Roberts procedure should be the first choice in sequential detection of business cycle turning points with leading indicators. At the same time parallel use of the CUSUM procedure could help detect fast and sharp turning points.

Appendix: Some Methodological Details and Graphs

Here we describe some details of our computations. As is seen from the definitions of the Shiryayev-Roberts and CUSUM procedures they can be used in detecting a change in one direction only. Therefore, in order to detect business cycle turning points we need to run each of them twice: the first time to detect recessions, and the second time to detect expansions. In detecting expansions we used the original recursive residual series, X_t, computed from (6). In detecting recessions we used series $-X_t$ (the original series multiplied by -1).

Since the EWMA procedure defined in this paper is two sided[3] it can detect changes in both directions. Therefore in this case there was no need to change the sign of the series – we used the original series of the recursive residuals. But the EWMA procedure defined in (4) performed poorly. This happened because of the well known problem (Lucas and Saccucci (1990)) that if at the time of a change in one side (positive, for example) the EWMA statistic is in the other side (negative) then there will be a delay in detection. One solution to this problem is to consider two EWMA procedures: one is designed to detect negative changes/recessions and the other one to detect positive changes/expansions. In the first case we reset the value of the EWMA statistics to zero whenever it reached or went above UCL, $L\sigma$. So that it helps the statistic go down to LCL more quickly whenever a negative change occurs. In the second case we reset the value of the EWMA statistics to zero whenever it reached or went below LCL, $-L\sigma$ for the same reasons. We could use slightly different but seemingly more reasonable resetting rule: in the first (second) case we could reset the EWMA statistic to zero whenever it becomes positive (negative). But there are two reasons for not considering this alternative. First, it would significantly reduce the ARL values making these EWMA procedures non comparable with the other two procedures. Second, this alternative could not improve the EWMA performance at all compared to the previous one while increasing the false alarm rate significantly.

Below we present the results in graphical form. In order to achieve a better view we present the graphs of the Shiryayev-Roberts and CUSUM procedures in the logarithms: we plot the logarithms of the control statistics[4] and therefore we use the logarithms of the thresholds as appropriate control limits.

[3] It has two control limits: the upper control limit(UCL), $L\sigma$, and the lower control limit (LCL), $-L\sigma$.
[4] In the case of the CUSUM procedure we plot $\log\{\max(1,\max_{1\leq k\leq n}(S_n - S_k))\}$.

Detecting recessions using CLI

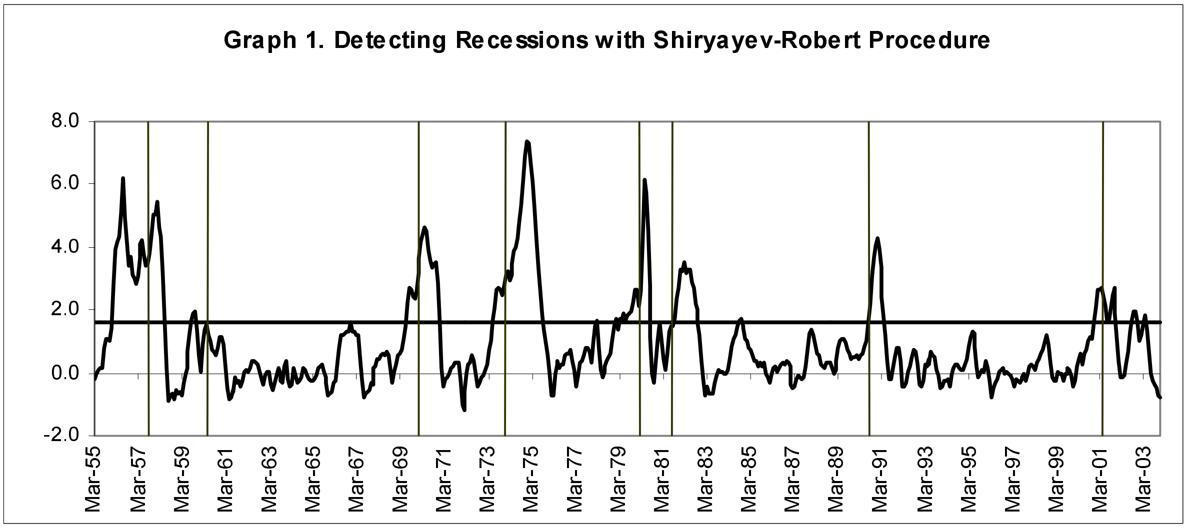

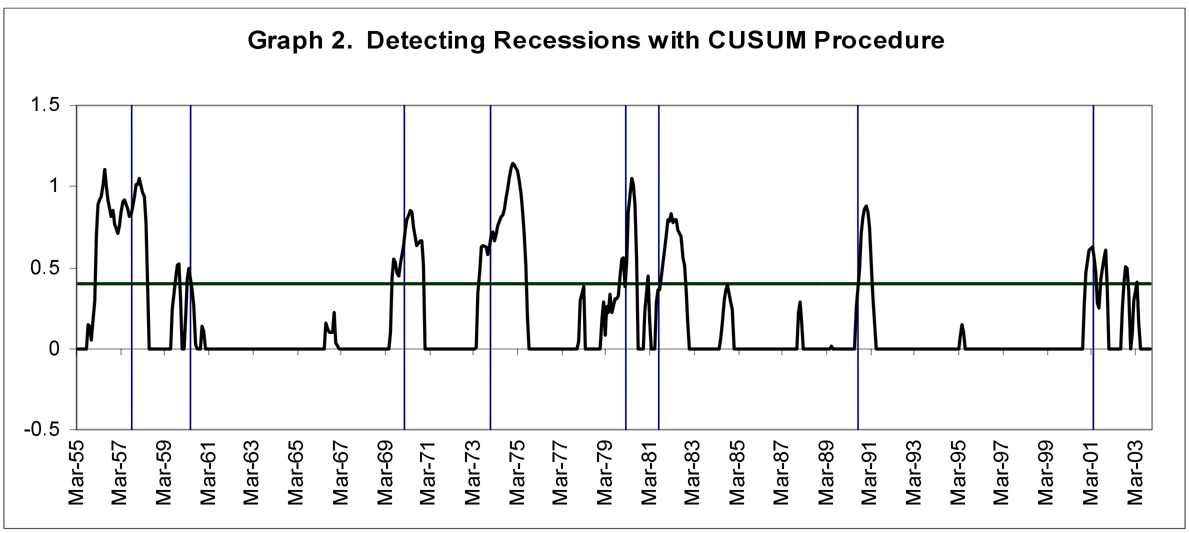

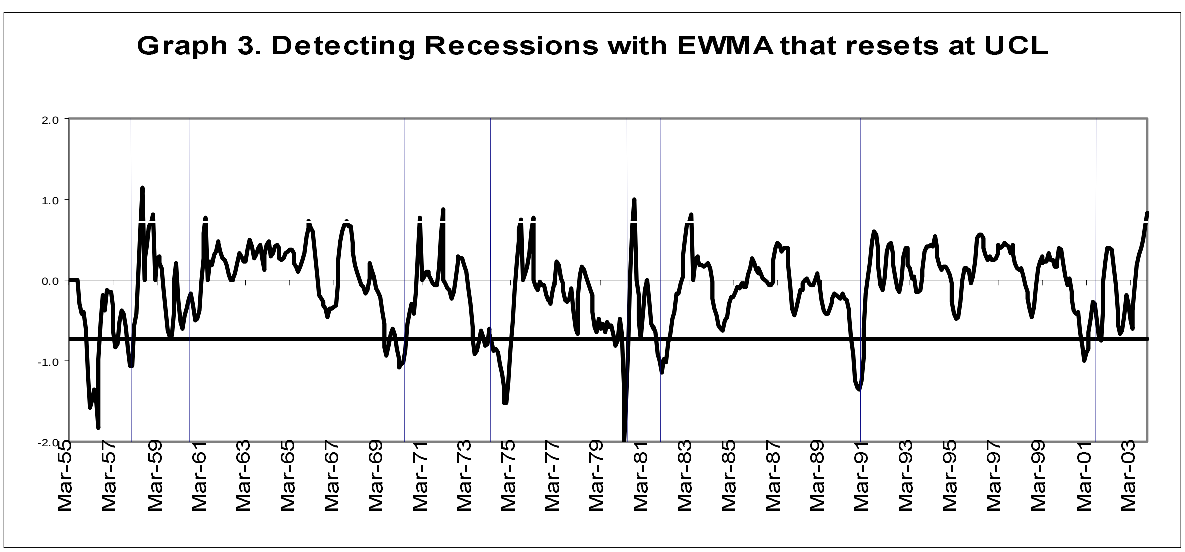

Detecting expansions using CLI

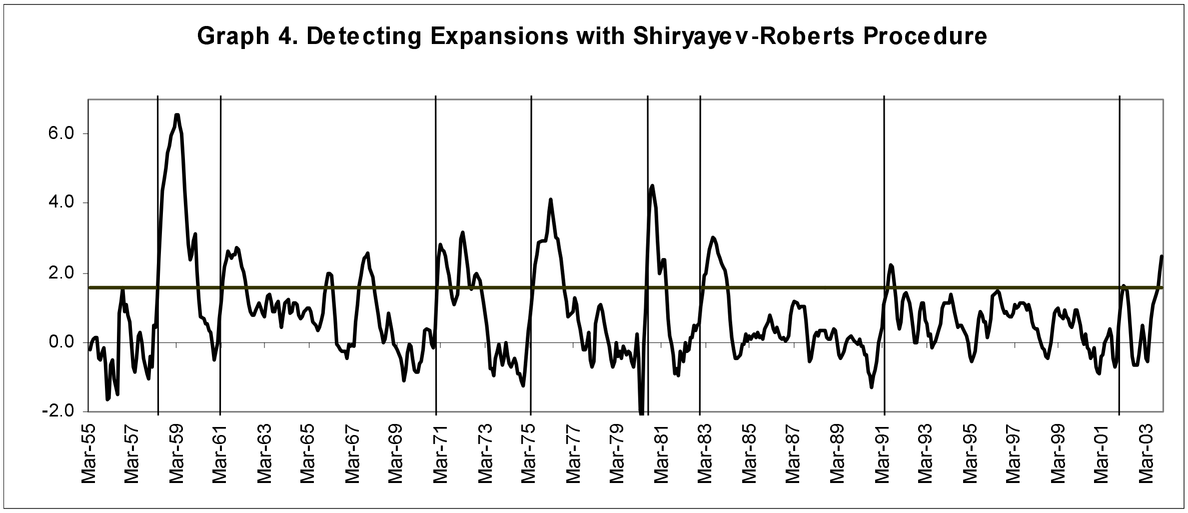

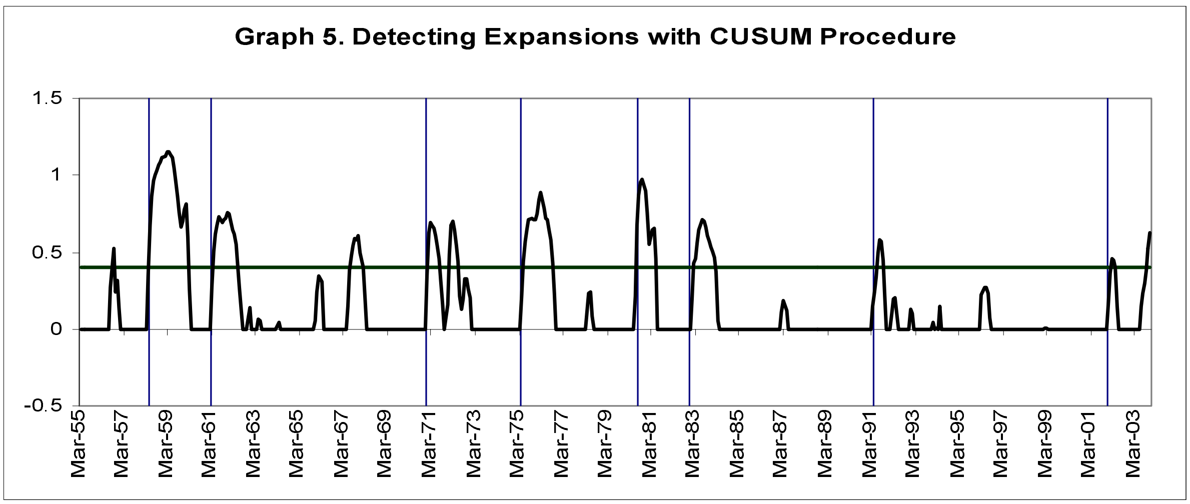

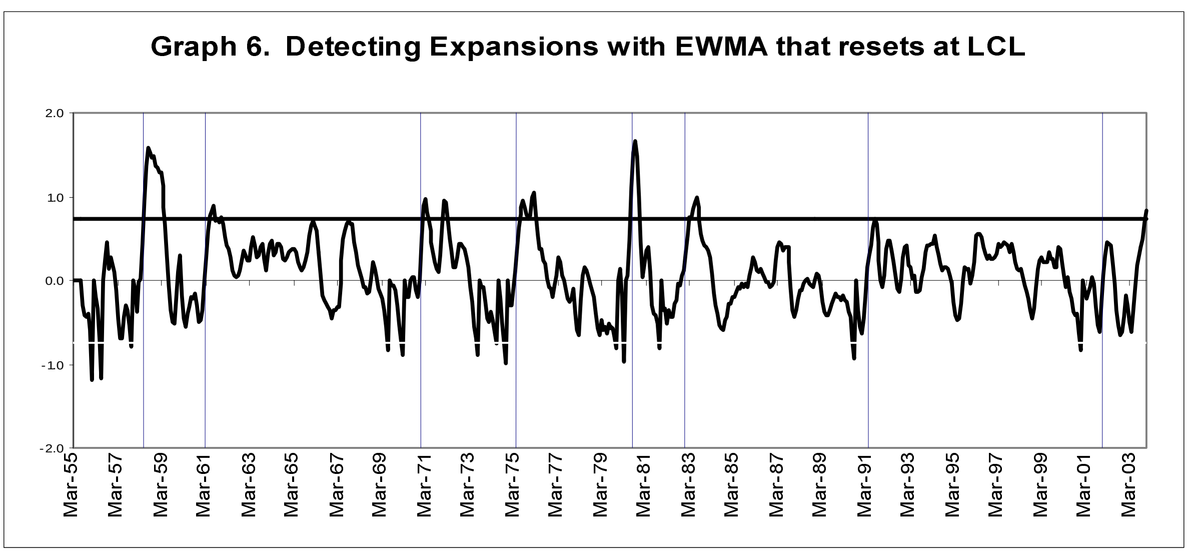

Detecting business cycle turning points using S&P Composite Index

Graph 7. Detecting Recessions with CUSUM Procedure using S&P Composit Index

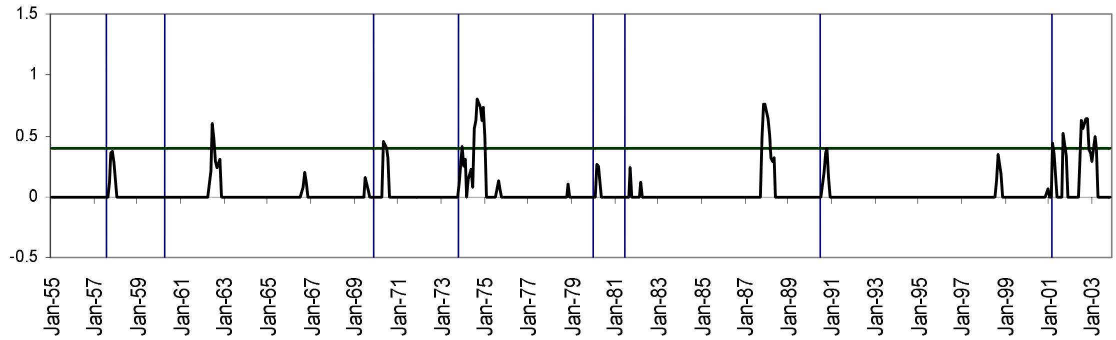

Graph 8. Detecting Expansions with CUSUM Procedure using S&P Composit Index

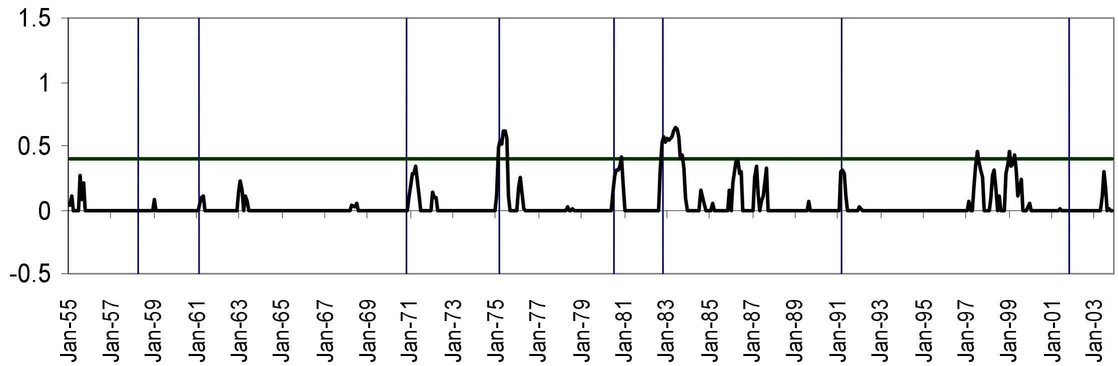

References

Andrews, Donald W.K., 1993, Tests for Parameter Instability and Structural Change with unknown change point, Econometrica, 61,821-856.

Bai, J. (1994) Least Square Estimation of a Shift in Linear Processes. Journal of Time Series Analysis, 15,453-472.

Bai, J. (1997) Estimation of a Change Point in Multiple Regression models. Review of Economics and Statistics, 79, 551-563.

Brown, R.L., Durbin, J., and Evans, J.M. (1975) Techniques for testing the constancy of regression relationships over time. Journal of Royal Statistical Society, B, v.37, 149-192.

Chib, S. (1998) Estimation and comparison of multiple change-point models. Journal of Econometrics, 86, 221-241.

Chow, G.C. (1960) Tests of equality between sets of coefficients in two linear regressions. Econometrica, 128, 591-605.

Cooley, T.F., and Prescott, E.C. (1973) Tests for an Adaptive Regression Model. Review of Economics and Statistics, 55, 248-256.

Cooley, T.F., and Prescott, E.C. (1976) Estimation in the Presence of Stochastic Parameter Variation. Econometrica, 44, 167-184.

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

Ergashev, B.A. (1991) Asymptotical solution of one problem of detecting a multi-alternative change point. Chapter 2 of unpublished dissertation "Some problems of Sequential Analysis related to the change point problem". Steklov Mathematical Institute, Moscow (in Russian).

Ergashev, B.A. (2003) On a CAPM monitoring based on the EWMA procedure. Working Paper. Presented at 9-th International Conference of the Society for Computational Economics and Finance, Seattle, July 11-13. Available online at:
http://logec.hhs.se/scripts/paperstat.pl?h=repec:sce:scecf3:283

Estrella, A., and Mishkin, F.S. (1998) Predicting U.S. Recessions: Financial Variables as Leading Indicators. Review of Economics and Statistics, 80, 45-61.

Fama, E.F., and French, K.R. (1993) Common risk factors in the returns on stocks and bonds. Journal of Financial Economics, 33, 3-56.

Guermat, C and Harris, R.D.F. (2001) Robust Conditional Variance Estimation and Value at Risk. Working paper 01/06, School of Business and Economics, University of Exeter.

Hansen B. E. (2000) Testing for Structural Change in Conditional Models. Journal of Econometrics. 15, 60-67.

Hansen, B.E.(2001) The new Econometrics of Structural Change: Dating Breaks in U.S. Labor Productivity. Journal of Economic Perspectives, 15, 117-128.

JP Morgan (1996), "RiskMetrics - Technical Document," Fourth Edition, New York

Knoth, S. (2002) Statistical Process Control, in Applied Quantitative Finance: theory and computational tools, Springer, pp. 237-258.

Lorden, G. (1971) Procedures for reacting to a change in distribution. Annals of Mathematical Statistics. V. 42, N. 6, 1897-1908.

Lucas, J. M. and Saccucci, M. S. (1990) Exponentially weighted moving average control schemes: properties and enhancements, with discussion. Technometrics, 32,1 -29.

Muller-Plantenberg, N. (2003) Current account reversals triggered by large exchange rate movements. Working paper, Economics Department, London School of Economics.

Neftci, S. N. (1982) Optimal Prediction of cyclical downturns, Journal of Economic Dynamics and Control 4, 225-241.

Neftci, S.N. (1984) Are Economic Time Series Asymmetric over the Business Cycle? Journal of Political Economy, 92, 307-328.
Novikov, A. A. (1990) On the first passage time of an autoregressive process over a level and an application to a "disorder" problem. The theory of probability and its applications, 35, 269-279.
Page, E.S. (1954) Continuous Inspection Schemes. Biometrika, 41, 100-115.
Pollak, M. (1985) Optimal detection of a change in distribution. Annals of Statistics,13, 206-227.
Pollak, M. and Siegmund, D. (1985) A diffusion process and its applications to detecting a change in the drift of Brownian motion. Biometrika, v. 72, 267-280.
Quandt, R. (1960) Tests of the Hypothesis that a Linear Regression Obeys Two separate Regimes. Journal of American Statistical Association, 55, 324-330.
Roberts S.W. (1966) A comparison of some control chart procedures. Technometrics, v. 8, 411-430.
Shiller, R.J. (2000). Irrational Exuberance. Princeton University Press. Princeton, NJ.
Shiryayev, A. N. (1963) On optimal method in earliest detection problems. Theory of Probability and its Applications, 8, 26-51.
Shiryayev, A. N. (1978) Optimal stopping rules, Applications of mathematics. Springer-Verlag, Berlin.
Siegmund, D. (1985) Sequential Analysis: Tests and Confidence Intervals. New York: Springer-Verlag, 272 p.
Stock J.H., and Watson, M.W. (1994) Evidence on Structural Instability in Macroeconomic Time Series Relations. NBER Working Paper No 164.
Zivot, E., and Andrews, D.W.K. (1992) Further evidence on the Great Crash, the oil-price shock, and unit-root hypothesis. Journal of Business and Economic Statistics, 10, 251-270.