DATA MINING FOR UNUSUAL MOVEMENTS IN TEMPORAL DATA

R. Douglas Martin[1] and Victor Yohai[2]

[1] Chief Scientist, Insightful Corporation and Professor of Statistics, Univ. of Washington

[2] Dept. of Mathematics, University of Buenos Aires, Buenos Aires, Argentina

Abstract

In data mining one is often confronted with a very large number of time series, e.g., thousands of stock prices over time, monthly spending on cellular services by large groups of consumers, or temporal records of WEB site visits. In virtually all such cases one finds that some fraction of the time series in question exhibit “unusual” movements of some kind, e.g., a temporally isolated “outlier” that differs radically from neighboring values, a “level shift” whereby the values of the series are all shifted up or down after a certain point in time. For example, one needs to know whether an unusual movement in a stock price has occurred, and whether it is an isolated outlier that might be a “true” value or an error in transmission, or a level shift that represents a locally permanent shift in value of the stock. In a similar vein, when wireless companies monitor billing for cellular services, it is imperative to accurately detect unusually large or small values in services used in a given month and to quickly determine whether it is a transient occurrence or the beginning of a new pattern of usage by the customer. In this paper we discuss some very useful approaches to detecting unusual movements using both simple and complex time series models, combined with robust statistical methods for time series. An important fact revealed in the discussion is that conventional non-robust statistical methods often fails to detect unusual movements. We provide a compelling example of this differentiated behavior of classical versus robust statistics with a simple application to the new Standard and Poors credit index. For this kind of time series, a simple random walk model suffices, and we show that use of conventional prediction error standard deviation estimates (“volatility” estimates) leads to an unusual movement detection statistic that fails to detect obviously significant movements. On the other hand a simple robustification of the volatility estimate leads to very effective detection of unusual movements, i.e., a detector with high statistical power. We also discuss more complex models that may sometimes be needed, e.g., regression models with correlated autoregression moving-average type errors, and Bayesian dynamic structural models. In the oral presentation of this paper we shall briefly discuss trade-offs between using more sophisticated models versus simpler models. The latter, when effective are obviously more attractive for data mining very large numbers of time series.

1. INTRODUCTION

One of the most important and relatively neglected problems in applications of data mining to temporal data, is the automatic detection of sudden unusual movements. This is a problem that has been studied fairly extensively in the statistical literature, where temporal data is referred to as *time series* data. The purpose of this paper is to introduce a few methods suggested by this literature that have considerable potential for use in data mining applications that involve possibly very large numbers of time series. In particular we will introduce the use of robust methods for handling isolated time series outliers, and subsequently indicate how these methods can be integrated with time series structural modeling to detect abrupt changes in the level of time series. Along the way, we will point out some of the relevant statistical literature on robustness and change detection for time series.

1.1 Types of Unusual Movements in Time Series

One can imagine an almost unlimited number of types of sudden unusual movements in temporal/time series data. However, the following basic types of "sudden" unusual movements stand out as both commonly occurring, and in need of automatic and rapid data mining detection methods:

- Isolated outliers
- Level shifts
- Shifts in variability
- Slope changes
- Changes in frequency

Of these we shall concentrate here on the first two, briefly touch on the third, and in the summary section point out some relevant literature on the last two.

Figure 1 below shows a time series of monthly tobacco and related sales in the UK from 1955 to 1960 (from West and Harrison, 1989). Aside from the overall pattern of growth in sales starting at the beginning of 1955 through 1958, and the subsequent slow decline in sales during 1959, the most striking features of the data are: the *isolated outlier* at the end of 1958, and the three distinct levels of sales during 1956, 1957 and 1958, with abrupt *level shifts* of sales occurring at the beginning of 1957 and 1958.

The two distinct sudden change types in Figure 1 occur frequently in temporal data. From a data mining perspective it is obviously of paramount importance to be able to detect such sudden unusual movements in temporal data, and to distinguish isolated outliers from level shifts. An isolated outlier could be due to a sudden temporally localized spending splurge by a customer, or one-time fraudulent use of another persons phone number to place high-cost international calls during a particular month, or the *return* (relative change in price) due to an isolated movement of a stock price to a new level, or simply an erroneous data point (e.g., a transmission error in stock price quote).

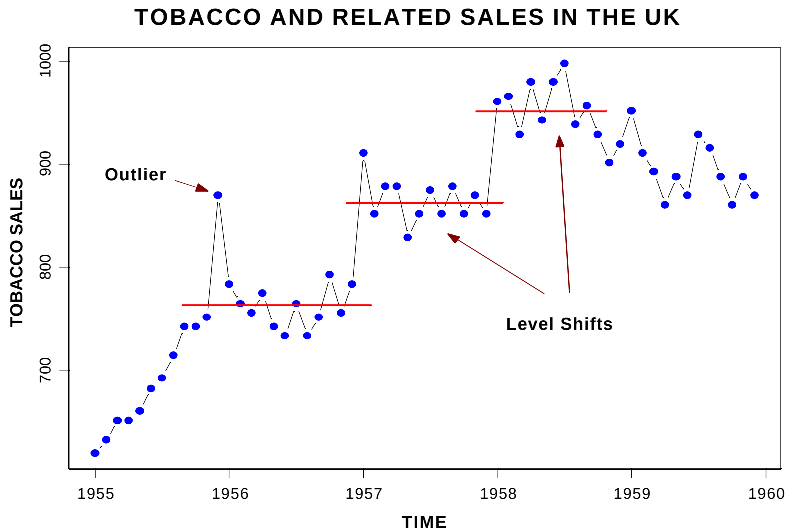

Figure 1. Time Series with Isolated Outlier and Two Level Shifts

A level shift on the other hand represent a change in value that endures for at least awhile, and represents a new *local mean* value for the temporal data. This might be caused be a stock price moving upward or downward to a new local level, or a sudden decrease in a customer's use of cellular roaming services after discovering the increased expense of roaming services, and so on.

Whatever the cause of sudden change, whether it be an isolated outlier or a level shift, there are obvious business payoffs in detecting such events as quickly as possible, and taking appropriate remedial action. For example, quick detection of an isolated outlier by the data integrity group in a financial data services organization can lead to quick verification as to whether to unusual stock price is a true movement in price or an erroneous recording or transmission error. And rapid detection of a cellular customer's decreased use of roaming services may lead the company to make a special offer to retain the customer's use of roaming services.

If one knows based on previous experience that only one of the above two sudden unusual movements can occur, either only isolated outliers occur or only level shifts occur, than it is much easier to build a data mining algorithm to rapidly detect and report such an occurrence. When both level shifts and outliers can occur, detecting the unusual movement is still relatively easy, but the problem of distinguishing which occurred and what subsequent action one should take is more challenging. For having detected an unusual movement at the time of occurrence of its occurrence, one typically can not know whether it is an isolated outlier or a level shift. Instead it is necessary to observe one or more data values subsequent to the unusual movement are needed in order to make a reliable decision.

The case where only isolated outliers occur is the simplest case, and it is one for which robust statistical methods for time series with outliers was invented. See for example, Martin and Denby (1979), Martin and Yohai (1985). The case where one has to detect level shifts only is somewhat more complicated, and requires either an ad hoc detection method or an approach based on statistical tests for change in a structural time series model for level shifts. Dealing with the general case where the series contains both outliers and level shifts can be dealt with in several ways, one of which is to generalize the approach for dealing with level shifts to include the possibility of outliers.

The remainder of this paper is organized as follows. In Section 2 we show how to construct a fairly simple robust method of detecting unusual movements in financial time series where the primary concern is with obtaining a robust local volatility estimate as a key ingredient of the detection method. In this context, a simple approach is possible because the primary concern is with outliers in the differenced series, and the efficacy of the method is demonstrated on the new S&P Credit Index. In Section 3, we discuss extensions of the method of Section 2 to more general volatility models and to sudden changes in volatility. In section 4 we introduce the use of a general class of robustly fitted time series models to detect both isolated outliers and level shifts, and to distinguish between these two types of sudden change. Section 5 provides some summary comments and pointers to some of the relevant time series literature.

2. ROBUST DETECTION OF UNUSUAL MOVEMENTS IN FINANCIAL TIMES SERIES

Figure 2 shows a time series of daily values of the Investment Grade S&P Credit Index in units of *basis points* 120 days during 1998, along with the first differenced Index series.

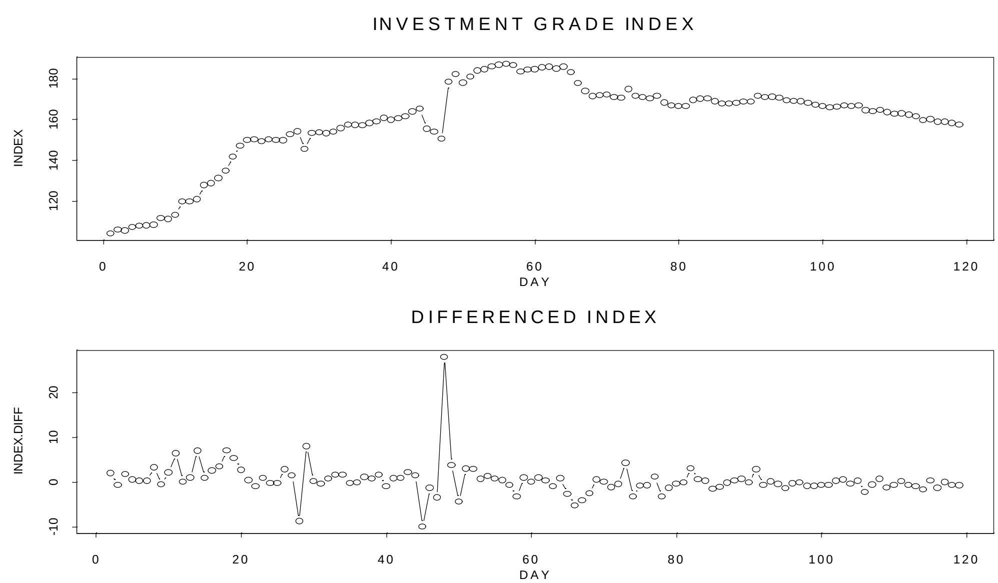

Figure 2. The S&P Investment Grade Credit Index and its First Difference Series

Aside from an overall upward trend for the first half of the series and slower downward trend for the second half, the most striking features of the Index are and isolated outlier that is locally (in time) negative, and a short negative level shift.

Notice that in the differenced Index series, the isolated outlier becomes a "doublet" outlier that first goes negative and then goes positive, and that the local level shift results in a negative isolated outlier followed by an isolated positive outlier in the differenced series. You also notice that the local *volatility*, i.e., the local standard deviation of the differenced Index series, varies slowly over time and is higher at the beginning of the series than at the end of the series. Note also, that the differenced Index series has no level shifts, only outliers.

Our focus here will be on detecting the unusual movements in the Index without regard to quickly deciding whether or not an isolated outlier or a level shift in the Index has occurred. It turns out that in order to do the key to success is the construction of a robust estimate of the volatility of the differenced Index series in the lower plot above, and the essence of the method involves one-sided interpolation (i.e., prediction) at the time of occurrences of the outliers.

We remark that the unusual movements behavior exhibited by the S&P Credit Index and its differenced series is a commonly occurring one with many financial time series such as stock prices, foreign exchange rates, etc., where one looks at *returns* instead of relative differences (returns are the relative change in price from one period to the next, i.e., the first difference normalized by the price at the beginning period). Thus the methodology described here is fairly generally applicable to a broad class of financial time series.

2.1 The Random Walk Approximation for the Credit Index Time Series

When you check for serial correlation (autocorrelation) in the differenced series you find out that there is little evidence of serial correlation. It is therefore reasonable to model the Index as a random walk

$$x_t = m + x_{t-1} + r_t \tag{1}$$

where m is the drift parameter and the random shocks r_t are zero mean white noise, i.e., zero mean and serially uncorrelated. Let s_t^2 be the variance of r_t. With this notation the *volatility* of x_t is s_t. The "local" mean of the differenced series, i.e., the mean over a moving window of say 20 days or so, does not often deviate from zero by more than a few basis points. Thus it is reasonable to use the approximation $m = 0$, in which case we have

$$r_t = \Delta x_t = x_t - x_{t-1}. \tag{2}$$

2.2 The Classical EWMA Volatility Estimate

With $m = 0$ we can use the first differenced series $r_t = \Delta x_t$, t = 2, …, n, to estimate the slowly varying volatility s_t^2. The best-known simple and reasonably effective way to do this is with the *exponentially-weighted moving average* (EWMA) estimate defined by the recursion

$$\hat{s}_{t+1}^2 = l \cdot \hat{s}_t^2 + (1 - l) \cdot \mathrm{r}_{t+1}^2, \quad \mathrm{t} = \mathrm{t_o}, \mathrm{L}, n \tag{3}$$

with initial condition $s_{t_o}^2$, and with $0 < l < 1$. If you take $t_o = -\infty$ and iterate the recursion backward a step at a time, you get

$$\hat{s}_t^2 = \sum_{i=-\infty}^{t} l^{t-i} \cdot r_i^2 .$$

which reveals why the estimate has the name it does.

In practice, one starts the recursion by computing the initial condition $s_{t_o}^2$ for (3) as the sample standard deviation of an initial fraction of the data.

2.3 An Unusual Movement Detection Method for Normally Distributed Data

An unusual movement of the index occurs at time t whenever the absolute value $|\Delta x_t| = |r_t|$ is "sufficiently" large. The natural yardstick for this measurement is the local volatility estimate $\hat{s}_t$. Therefore one might use the absolute value of the following *unusual movement test* statistic

$$UMT_t = \frac{r_t}{\hat{s}_t} . \tag{4}$$

which is a natural analog of the classical statistics t-test. If the r_t in (8) are normally distributed and the value of l used in computing $\hat{s}_t$ is in the vicinity of .9 or larger (so that $\hat{s}_t$ is based on sufficient weighted averaging of past data to have little variability), then UMT_t will be approximately normally distributed. In that case we can characterize an unusual movement by using the standard normal distribution to set a threshold *thresh* such that the probability that $|UMT_t|$ exceeds *thresh* is any desired user-specified probability p. Table 1 below gives a few values to choose from.

p	thresh
.01	2.58
.005	2.81
.001	3.29
.0005	3.48

Table 1. Normal Distribution Thresholds for Various p-Values

We propose using the value *thresh* = 3.29 for which $p = .001$. The interpretation then is that in situations where there are in fact no unusual movements of the index, one will see a value of $|UMT_t|$ larger than 3.29 only one times in a thousand. If one wants to detect movements that occur even less (more) frequently when in fact no unusual movement occurs, then a larger (smaller) threshold can be used.

In order to check the normality assumption underlying the use of UMT_t, Figure 3 below provides a normal QQ-plot of 98 UMT_t values, where $\hat{s}_t$ in (3) has been computed

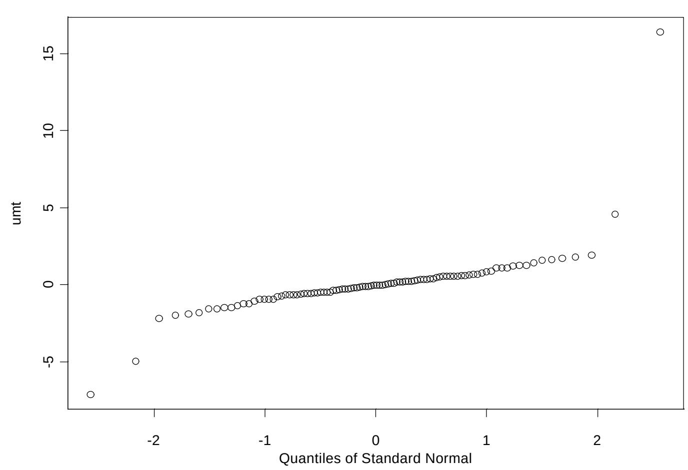

Figure 3. Normal QQ-Plot of UMT Values for the S&P Credit Index

using a value $l = .9$ and with $t_o = 20$ in the initial condition estimate $s_{t_o}^{2}$. The linearity of the plot, except for the few outliers at each end that correspond to the outliers generated by the unusual movements in the Index, is strong confirmation that we are justified in assuming that the UMT_t are normally distributed when there are no unusual movements in the series. However, this is not enough to insure that UMT_t provides a powerful detector of unusual movements, as the next subsection shows.

2.4 The Failure of the Normal Distribution Theory UMT in the Presence of Outliers

If we apply the above classical EWMA volatility estimate to the differenced index time series in the lower plot of Figure 1, we get the result shown in the top plot of Figure 4 below.

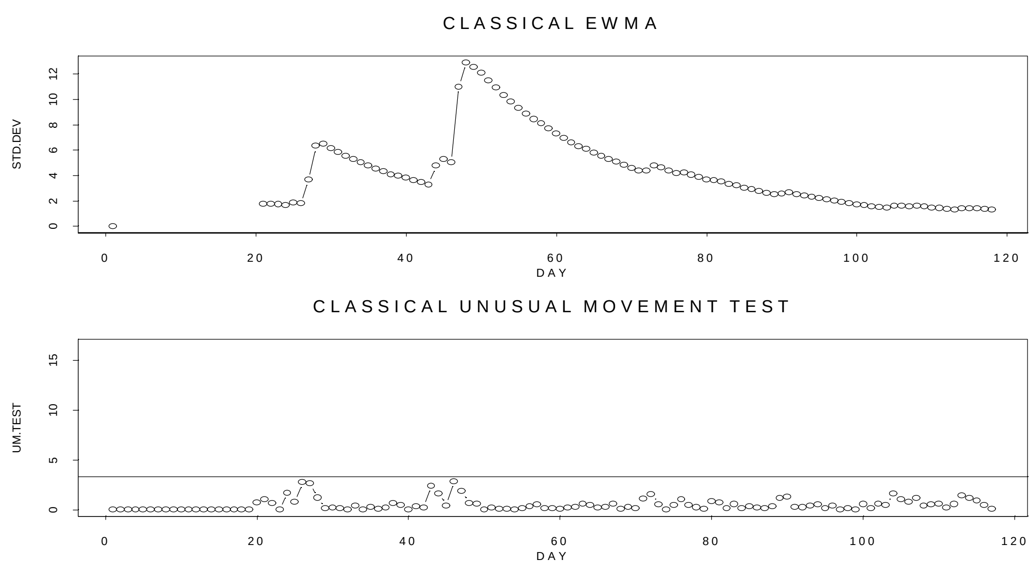

Figure 4. The Classical EWMA Volatility Estimate and Classic UMT Behavior

The bottom plot of Figure 4 shows the values of the test statistic $|UMT_t|$, with a horizontal line drawn at the value *thresh* = 3.29. The first 19 values are set at zero in each plot because the first value of $\hat{s}_t$ is available only starting at $t_o = 20$. Since none of the values of $|UMT_t|$ exceed the value 3.29, no unusual movement is detected! But it is qualitatively quite clear from Figure 1 that there are at least three very unusual movements that should be so declared by a good unusual movement detection method.

So what is wrong with our proposed test statistic? The problem is evident in the top plot of Figure 4. The estimate $\hat{s}_t$ takes several unusually large positive jumps, particularly in the vicinity of day 27 and in the vicinity of day 48. The reason for the jumps is that the estimate $\hat{s}_t$ is not at all *robust* in that it is highly influenced by the outliers in the differenced Index time series in the lower plot of Figure 1. The way around this problem is to create a robust version of the EWMA estimate of volatility UMT_t, and use this robust estimate of volatility as the denominator in the denominator of the UMT_t statistic (8).

2.5 The Robust EWMA Volatility Estimate and Robust UMT

Before describing how to robustify the classical EWMA volatility estimate to protect against the influence of outliers, we show the result of using it and the resulting robust UMT_t in Figure 5 below. Now the estimate $\hat{s}_t$ no longer grossly overestimates the volatility in the time periods following the occurrence of the outliers in the differenced

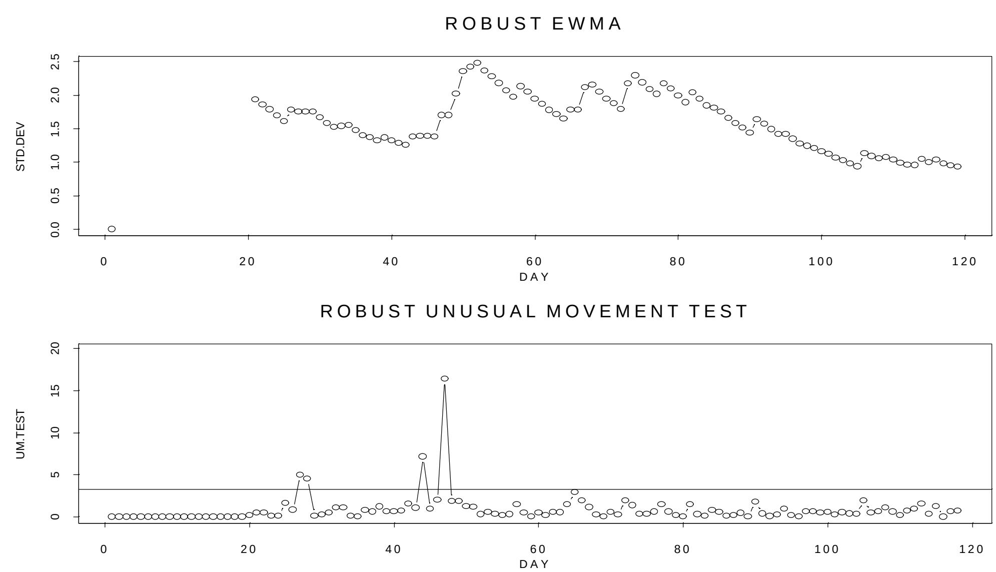

Figure 5. The Robust EWMA and Robust UMT

Index series, and in fact rather nicely estimates the slowly varying volatility of the differenced Index series. Correspondingly, the denominator of UMT_t is now much smaller at the occurrence times of the outliers in the differenced Index, resulting in values of UMT_t substantially larger than the threshold and thereby detecting the unusual movements.

We note that is you keep track of the signs of the outliers in the differenced Index, you can actually decide on the basis of the above plot that the first movement was an isolated outlier while the second movement was probably a level shift. This provides one relatively inexpensive way to discriminate between isolated outliers and level shifts in the original Index series. A more refined method of accomplishing this is described in Section 4.

The robustified EWMA volatility estimate is computed as follows. Set a threshold value *a*, and compute the robust volatility estimate using the following recursion for the robust variance (the square of the robust volatility estimate): For $t \geq t_o$, let

$$\begin{aligned} \hat{s}_{t+1}^2 &= l \cdot \hat{s}_t^2 + (1-l) \cdot r_{t+1}^2, && \text{if } \left|r_{t+1}\right| \leq a \cdot \hat{s}_t \\ &= \hat{s}_t^2, && \text{if } \left|r_{t+1}\right| > a \cdot \hat{s}_t \end{aligned} \tag{5}$$

using a robust scale estimate as the initial estimate $\hat{s}_{t_o}$. For robust initial scale estimate we used the median absolute deviation about the median, which is known to have highly desirable robustness properties (Martin and Zamar, 1989, 1993). In Figure 4, we have used *a* = 2.5 based on experience in using robust estimates of scale. It does not seem to make too much difference if you use values of *a* as low as 2 or as high as 3.

It is interesting to note that the robust EWAM algorithm (5) is equivalent to setting

$$r_{t+1}^2 = \hat{s}_t^2, \qquad \text{when } \left|r_{t+1}\right| > a \cdot \hat{s}_t \tag{6}$$

in (3), and leaving r_{t+1}^2 unaltered otherwise. That is replace r_{t+1}^2 with the predicted value $\hat{s}_t$ when r_{t+1}^2 is too large, and leave it alone otherwise.

3. OTHER ROBUST VOLATILITY ESTIMATES

The key to the above robust unusual movement test is the robust volatility estimate. The EWMA approach is a simple and often used model for volatility estimation, but there are many other important volatility estimation models that are sometimes preferred, such as the *autoregressive conditionally heteroscedastic* (ARCH) models of Engle (xxxx), and the *generalized autoregressive conditionally heteroscedastic* (GARCH) models introduced by Bollerslev (xxxx) and extended by many other authors. These models may be robustified in a spirit similar to that we used above. For example, robust ARCH models are described by Yohai xxxxx

4. DETECTING OUTLIERS AND LEVEL SHIFTS USING REGRESSION MODELS WITH ARIMA ERRORS

In the statistical time series literature one of the most frequently occurring general class of models used to fit and forecast time series data is the class of *linear regression* models, with correlated errors described by an *autoregression integrated moving-average* (ARIMA) model. These so-called *REGARIMA* models are specified as follows:

$$\boldsymbol{y_t} = \boldsymbol{x}_t' b + e_t \qquad t = 1, \mathrm{L}, T$$

where $\boldsymbol{y_t}$ is the observed time series, $\boldsymbol{x_t}$ is a p-dimensional column vector of predictor variables, b is a p-dimensional column vector of unknown regression coefficients. The error term e_t follows the ARIMA process:

$$\Phi(\boldsymbol{B})(1-\boldsymbol{B})^{d}(1-\boldsymbol{B}^{s})^{D} \cdot e_t = \Theta(\boldsymbol{B}) \cdot \boldsymbol{u_t}$$

where B is the lag (or backshift) operator, d the number of regular differences, D the number of seasonal differences, s the seasonality frequency, $\boldsymbol{u_t}$ is an independent and identically distributed (white noise) process with mean zero and variance $s_{\boldsymbol{u}}^{2}$, and where

$$\Phi(\boldsymbol{B}) = 1 - f_1\boldsymbol{B} - f_2\boldsymbol{B}^2 - \quad \mathrm{L} \quad - f_p\boldsymbol{B}^p$$

is a *stationary autoregression operator of order p* called an AR(p) operator, and

$$\Theta(\boldsymbol{B}) = 1 - q_1\boldsymbol{B} - q_2\boldsymbol{B}^2 - \quad L \quad - q_{\boldsymbol{q}}\boldsymbol{B}^{\boldsymbol{q}}$$

is a *stationary moving-average operator of order q*, called an MA(q) operator.

The usual approach to fitting the above regression with ARIMA errors to time series data is to assume that the $\boldsymbol{u_t}$ have a Gaussian ("normal") distribution, and use the method of maximum likelihood to estimate all the parameters in the model.

This classical statistical time series approach is not very useful for detecting outliers and level shifts in temporal data. However, a simple modification of the above model to account for isolated outliers and level shifts. Then use of a robust filtering method combined with robust test statistics and a robustified version of the the Gaussian maximum likelihood estimate results in a highly effective method. The modification of the model accounts for the possibility of isolated outliers and level shifts as follows.

4.1 A Modified REGARIMA Model to Account for Outliers and Level Shifts

The above REGARIMA model is modified as follows to account for outliers and level shifts.

Isolated Outliers

At time $\boldsymbol{t_o}$ the observation $\boldsymbol{y_t}$ in the above model is altered to

$$\boldsymbol{y}^*_{\boldsymbol{t_o}} = \boldsymbol{y}_{\boldsymbol{t_o}} + \boldsymbol{c}$$

where *c* is an arbitrary constant, whose effect is restricted to time $\boldsymbol{t_o}$.

Level Shifts

A level shift occurs at time $\boldsymbol{t_o}$ if

$$\boldsymbol{y}^*_{\boldsymbol{t}} = \boldsymbol{y}_{\boldsymbol{t_o}} + \boldsymbol{c} \qquad \textit{for all } t \geq t_o.$$

Note that if the series has an isolated level shift at time $\boldsymbol{t_o}$, the differenced series $\boldsymbol{y}^*_{\boldsymbol{t}} - \boldsymbol{y}^*_{\boldsymbol{t}-1}$ has an isolated outlier at time $\boldsymbol{t_o}$.

4.2 The Robust Filtering Algorithm

A robust filtering method is used to get estimates $\hat{e}_{t|t}$ of e_t that are not much influenced by previous outliers or level shifts. For the case of an AR(1) model the robust filter has the form

$$\hat{e}_{t|t} = \hat{e}_{t|t-1} + w_t \cdot (e_t - \hat{e}_{t|t-1})$$

where $\hat{e}_{t|t-1}$ is the one-step ahead prediction of e_t and

$$w_t = w\left(\frac{e_t - \hat{e}_{t|t-1}}{a \cdot \hat{s}_{t|t-1}}\right)$$

where w is an even and non-increasing weight function that goes to zero for sufficiently large argument, $\hat{s}_{t|t-1}$ is the one-step ahead prediction error variance and a is a tuning constant. This robust filter operates in the same qualitative way as the one we used to obtain robust volatility estimates in Section 2: when the prediction error is sufficiently large relative to the prediction error variance, the correction term in the above filter has value zero and the estimate $\hat{e}_{t|t}$ is the predicted value $\hat{e}_{t|t-1}$.

The form of the robust filter for general ARIMA models can be found in Martin, Samarov and Vandaele (1983).

4.3 The Robust Test Statistics

We compute the following unusual movement statistic:

$$T = \max_{t_0} \max\left\{ T_{t_o,O}, T_{t_o,LS} \right\}$$

where $T_{t_o,O}$ and $T_{t_o,LS}$ are robust test statistics for detecting isolated outliers and level shifts respectively. These test statistics have the same general robust t-statistic type form

$$T_{t_o,O/LS} = \frac{|\hat{c}|}{\hat{V}(\hat{c})^{1/2}}$$

where $\hat{c}$ is a robust estimate of the isolated outlier size or level shift size based on the residuals (prediction errors) obtained with the robust filter, and $\hat{V}(\hat{c})$ is a robust estimate of the variance of $\hat{c}$. The computed value of $\hat{c}$ depends upon which event is being tested. Details concerning these type of tests for outliers and level shifts (computed in a non-robust way) were introduced a long time ago by Fox (1972).

If T is larger than an appropriately chosen threshold x, then one declares that either an outlier or a level shift has occurred, and that the time t_o of occurrence and whether an outlier or a level shift has occurred are those where the double maximum above occurs. For recommended values of x, see Bianco, Garcia Ben, Martinez and Yohai (1997).

4.4 The Robustified Likelihood

Using the robust filter to obtain cleaned error predictor values $\hat{e}_{t|t-1}$ and corresponding prediction error variance estimates $\hat{s}_{t|t-1}$, one can use the standardized prediction residuals

$$r_t = \frac{e_t - \hat{e}_{t|t-1}}{\hat{s}_{t|t-1}}$$

in a robustified form of the Gaussian maximum likelihood to obtain robust parameter estimates for the REGARIMA model (upon which the above prediction errors depend), via an iterative nonlinear optimization.

The key aspect of the robustified likelihood is the replacement of the sum-of-squares of the error prediction errors $e_t - \hat{e}_{t|t-1}$ by a robust scale estimate. Details may be found in Bianco, Garcia Ben, Martinez and Yohai (1997).

An early version of the above procedure for the special case of autoregression integrated time series (i.e., with no moving average component), was introduced by Martin and Yohai (1996). Non-robust versions of the approach described here were introduced by Chang, Tiao and Chen (1988) and by Tsay (1988). Their methods are less reliable because they do not incorporate robust methods as an integral part of the estimation and testing process.

4.5 Applications Examples

Several examples of applying the above method of detecting outliers and level shifts were presented in Bianco, Garcia Ben, Martinez and Yohai (1997) and in Martin and Yohai (1996), as well as in the S+ChangeDetection module implemented in S-PLUS. One or two of these will be provided in the oral presentation of this paper.

5. DYNAMIC BAYESIAN MODELING METHODS

Another approach to detecting outliers and level shifts (and other changes such as slope changes, variance changes, etc.) is the use of dynamic Bayesian time series structural models as extensively described in West, Harrison and Stevens (1989). An attractive generalization of their technique was introduced by Bruce and Martin (1993), who compute a parsimonious history tree of posterior probabilities of the occurrence of outliers and level shifts. Application of this method to the Tobacco Sales example of Section 1 will be described in the oral presentation of this paper.

REFERENCES

Bianco, A., Garcia Ben, M., Martinez, E. and Yohai, V. (1996). "Robust procedure for regression models with ARIMA errors", *COMPSTAT '96 Proceedings of Computational Statistics, 27-38*. Physica-Verlag.

Bollerslev, T. (1986). "Generalized autoregressive conditional heteroscedasticity", *Jour. of Econometrics*, **31**, 307-27.

Bruce, A. G. and Martin, R. D. (1993). "Tree based robust Bayesian estimation of time series structural models", in *Proceedings of the 25th Symposium on the Interface.*

Denby, L. and Martin, R. D. (1979). "Robust estimation of the autoregression parameter", *Journal of the American Statistical Assoc*., **74**, 140-146.

Engle, R. F. (1982). "Autoregressive conditional heteroscedasticity with estimates of the variance of UK inflation", *Econometrica*, **50**, 987-1008.

Fox, A. J. (1972). "Outliers in time series", *Jour. of the Royal Statist. Soc*., **B**, **34**, 350-363.

Chang, I., Tiao, G.C., and Chen, C. (1988). "Estimation of time series in the presence of outliers", *Technometrics*, **30**, No. 2, 193-204.

Martin, R. D., Samarov, A. and Vandaele, W. (1983). "Robust methods for ARIMA models", in Applied Time Series Analysis of Economic Data, edited by E. Zellner.

Martin, R. D. and Yohai, V. J. (1985). "Robustness in time series and estimating ARMA models", invited paper for *Handbook of Statistics, Volume V: Time Series in the Time Domain*, edited by Hannan, Krishnaiah and Rao, North-Holland.

Martin, R. D. and Yohai, V. J. (1996). "Highly robust estimation of autoregression integrated time series models", *Publicaciones Previas No. 89, Facultad de Ciencias Exactas y Naturales*, Universidad de Buenos Aires.

Muler, N. and Yohai, V. J. (2001). "Robust estimates for ARCH processes", to appear in *Jour. of Time Series Analysis*.

S+ChangeDetection User's-PLUS 2000, preprint, Insightful Corporation.

Tsay, R. S. (1988). "Outliers, level shifts and variance changes in time series", *Jour. of Forecasting*, **7**, 1-20.

West, M. and Harrison, J. (1989). *Bayesian Forecasting and Dynamic Models*, Springer-Verlag.