Mercury Marine & Automatic Forecasting Systems

Data Cleansing and Automatic Procedures

David Reilly Sr. Vice President
Ranga Katti Director of Software Intelligence

MERCURY
"1 On The Water"

Mercury Marine

- Mercury Marine's forecasting process needed an improvement from basic approaches that did not provide accurate forecasts.

- We realized that our forecasts were a function of the software and its methodology. We did a survey of practices and we selected Autobox as a way to get things back in control

- We undertook an effort to track accuracies in a more rigorous way to understand our service levels

Researchers Lead The Way ! Some Developers Follow !

- Ruey Tsay “Outliers, Level Shifts, and Variance Changes in Time Series” *Journal of Forecasting* 1988, *7*(1), pp. 1.
- Ruey Tsay “Time Series Model Specification in the Presence of Outliers” *Journal of the American Statistical Association*, 1986, *81*(393), pp. 132-41
- Bell, W. (1983). "A Computer Program for Detecting Outliers in Time Series," in American Statistical Association 1983 Proceedings of the Business Economic Statistics Section, Toronto, pp. 624-639.
- Chang, I., and Tiao, G.C. (1983). "Estimation of Time Series Parameters in the Presence of Outliers," Technical Report #8, Statistics Research Center, Graduate School of Business, University of Chicago, Chicago.
- Chen, C. and G. C. Tiao (1990) "Random Level Shift Time Series Models, ARIMA approximations, and Level Shift Detection" The Journal of Business and Economic Statistics , January, 1990, p.81-96.
- Box, G.E.P., and Tiao, G. (1975). "Intervention Analysis with Applications to Economic and Environmental Problems," Journal of the American Statistical Association, Vol 70, pp. 70-79.

Data Cleansing

What is unusual?

- We are asked as children “What doesn’t belong?” We build a “model” in our head as to what is usual and what is unusual.

A	B
C	71

- Early researchers thought that all unusual values can be detected when they occur outside some pre-set range such as +- 3 sigma around the mean. This is only true when the expected value is equal to the mean and the values are uncorrelated with constant variance. Typically, data was plotted and a visual review was to identify anomalies

- Typically, statistics like the standard deviation were calculated and 3 sigma bands were put around the mean to identify outliers. The reality is that the standard deviation that was calculated is skewed upwards by the outlier itself so this approach is not reliable. An assumed model was used in this process where the data was subtracted by the mean to get residuals. Who is to say that the mean is in fact the correct model for the data?

MERCURY MARINE

What is unusual?

- When data needs to be cleansed this suggests that we have omitted an important variable in the modeling process. This omitted deterministic variable may be either known to us or unknown to us. Detecting this phenomenon often leads directly to “hypothesis generation” where data suggests theory, such as the need for an omitted event.

- Care must be taken not to falsely identify anomalies that are systematic such as a seasonal pulse variable.

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

What is unusual?

- We see a big outlier, but what about the pattern near the end?

- Do we remove/fix those also? Do we set them to be an average of the previous and next data points? Or do we identify those as “seasonal pulses” and include them as causal variables in the model so that they can be forecasted?

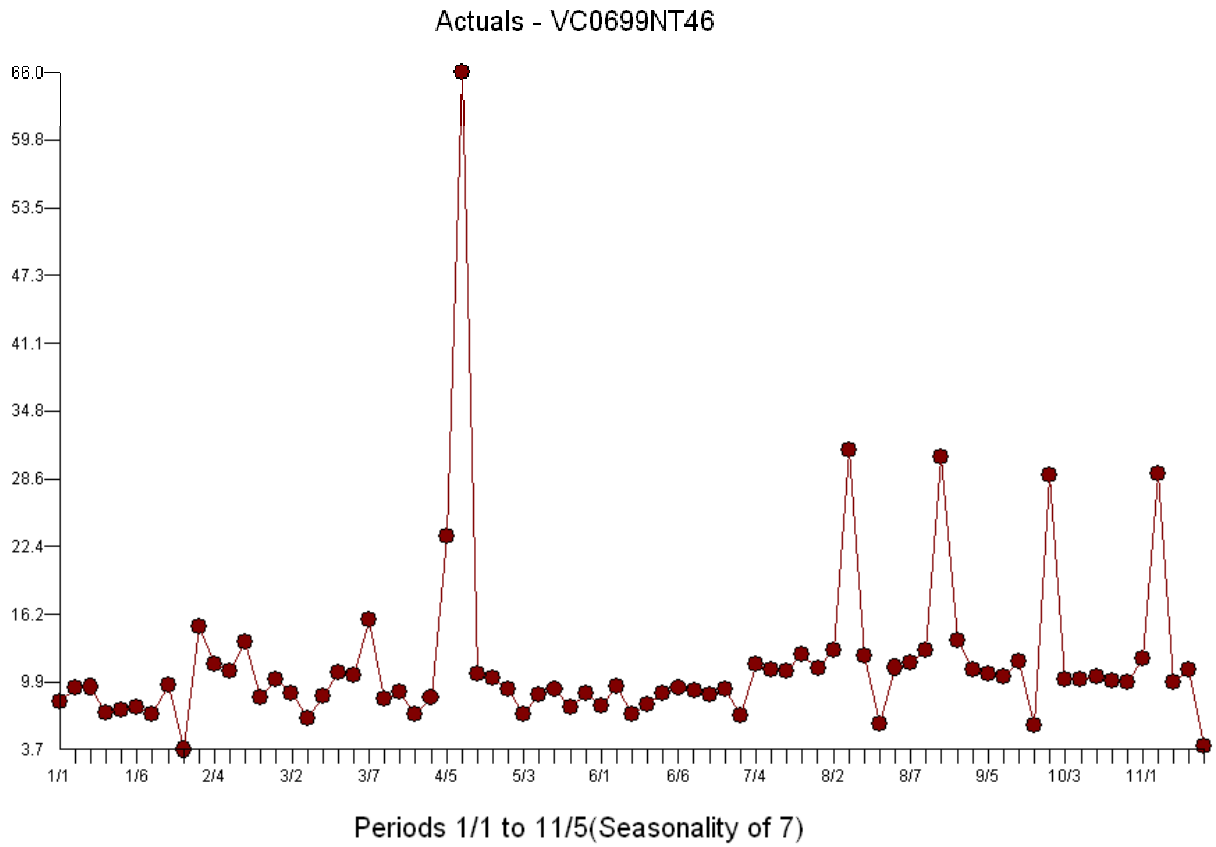

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

What is unusual?

- There are some outliers

- There is a seasonal pulse that begins in February near the end. If you don't account for this then the forecasts will use all Februaries to forecast and the forecast will be too low

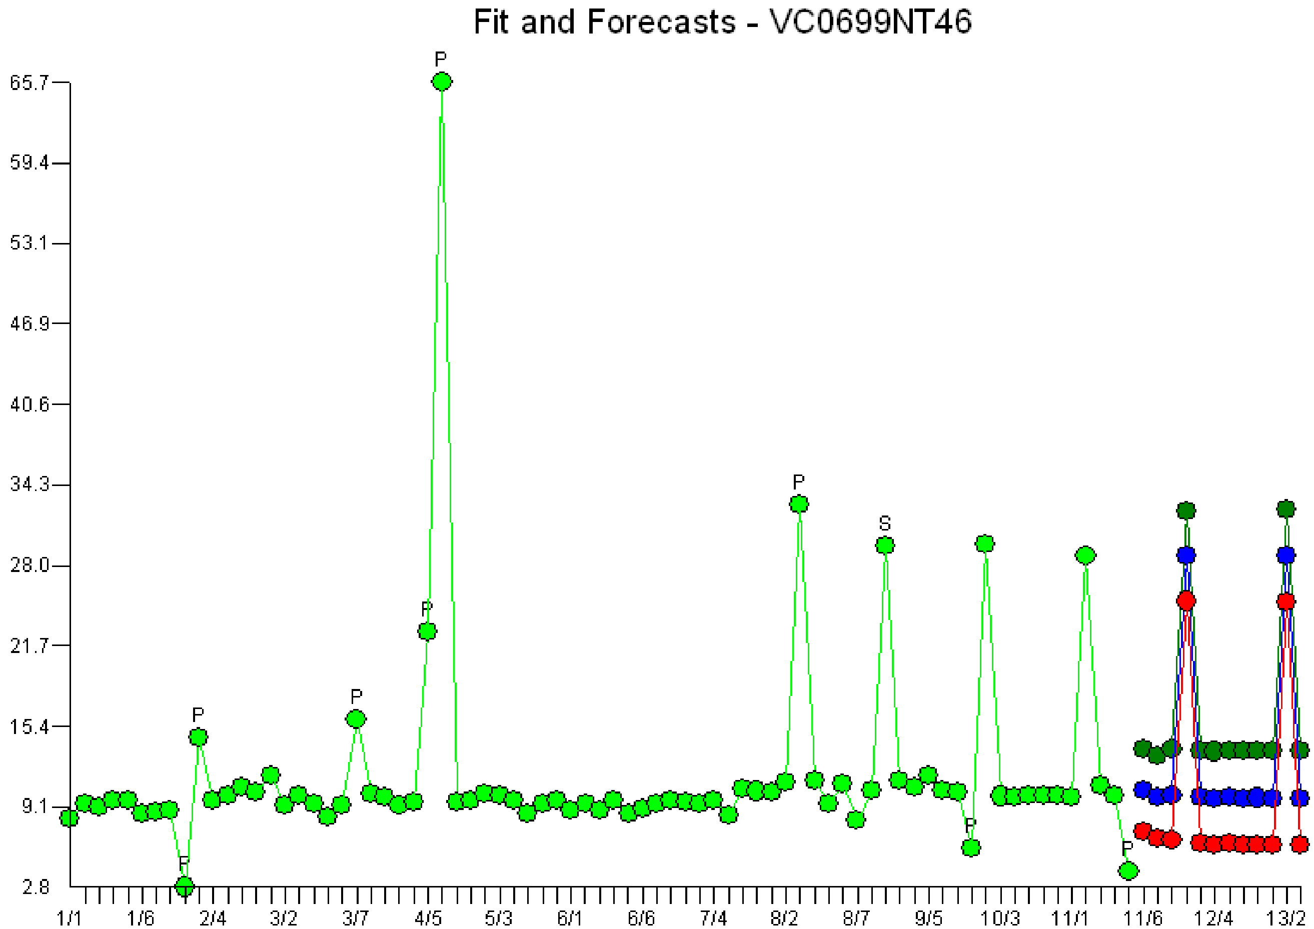

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Was it a Causal Model Issue all along?

- **We realize that we shouldn't be data cleansing at all. We should be adding causal information to the process. The culprit was that there was a buy one get one free (BOGOF) promotion that caused the change in demand.**

A '1' where there is a promotion and a '0' where there is no promotion

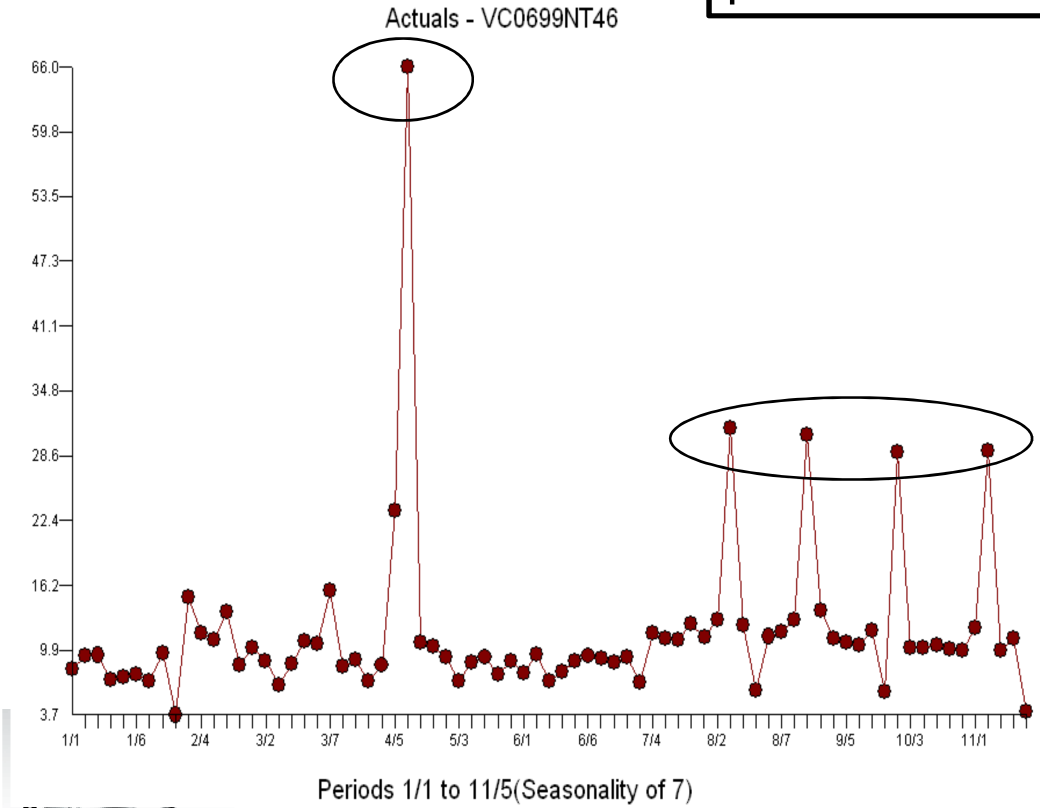

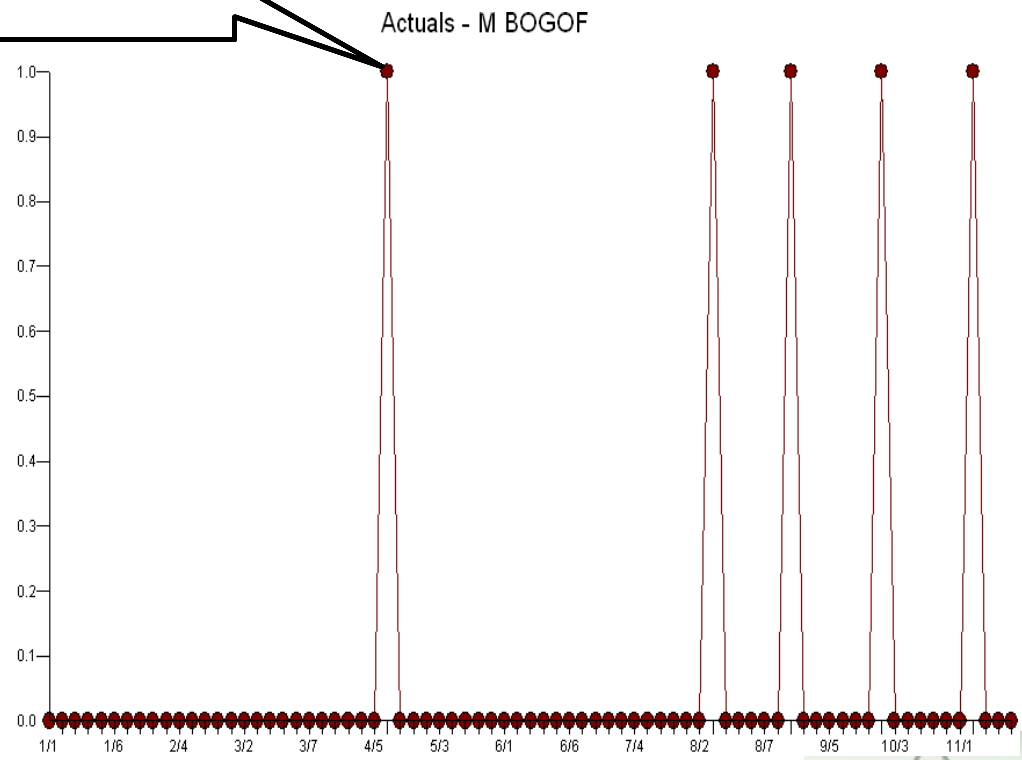

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

What is unusual?

- **This is U.S. enplanement data for 1997-2007. Note the decrease starting 9/11/01. How do you cleanse this dataset? There a few interventions in the fall of 2001 and when corrected for along with some new seasonal pulses, you are good to go.**

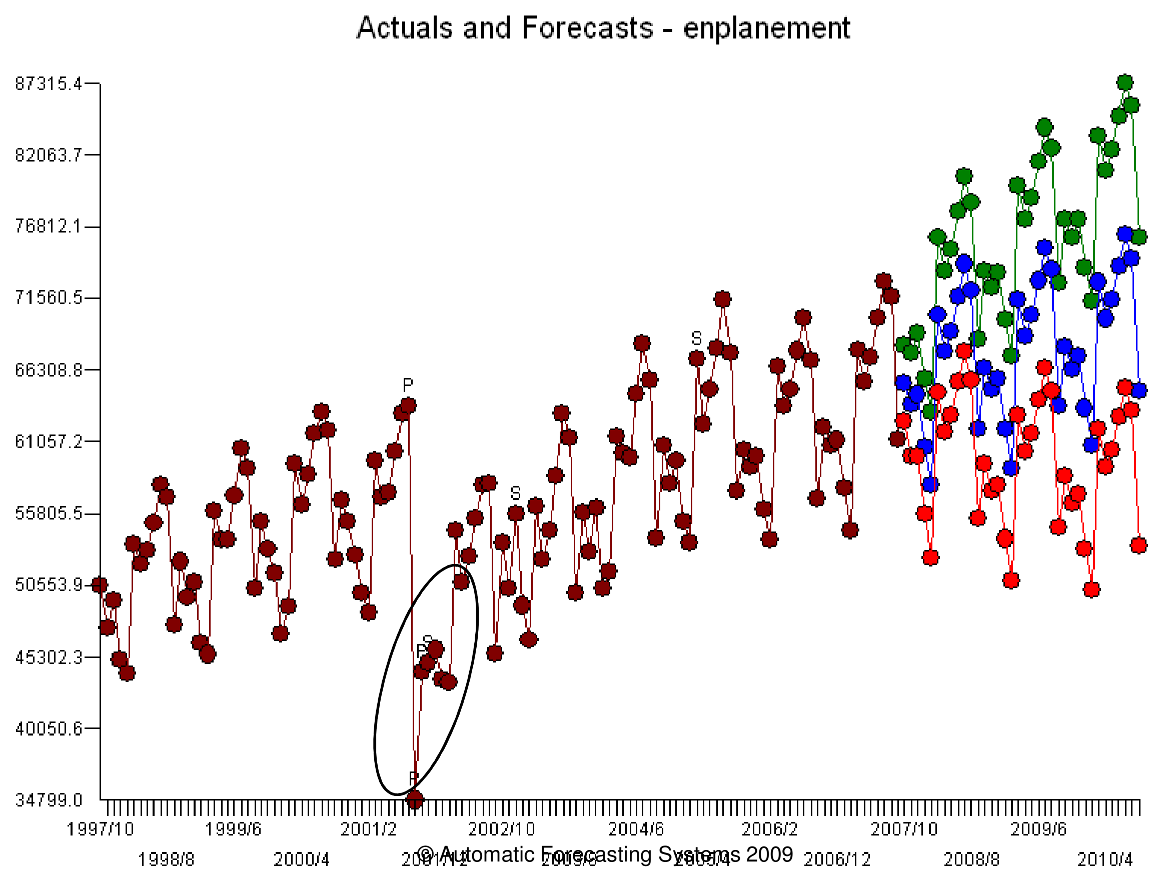

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

What is unusual?

- This is a example where the weekends have high sales. The last Saturday has a low value. Is this an "unusual value"? Yes, but how to identify and account for it. It is an <u>inlier</u> and the remedy is to "tweak" or adjust the observed value to ensure parameter optimization.

- If this value is not accounted for the model parameters and forecast will be affected

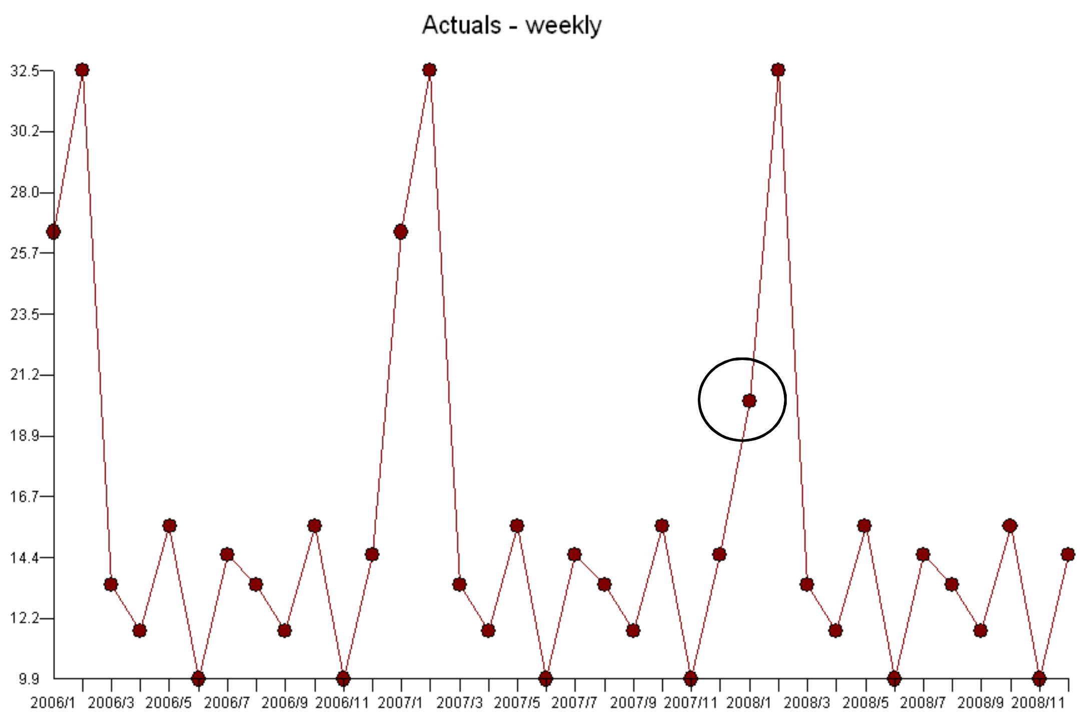

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Here is an outlier, right?

- **This value is <u>not</u> an unusual data point**

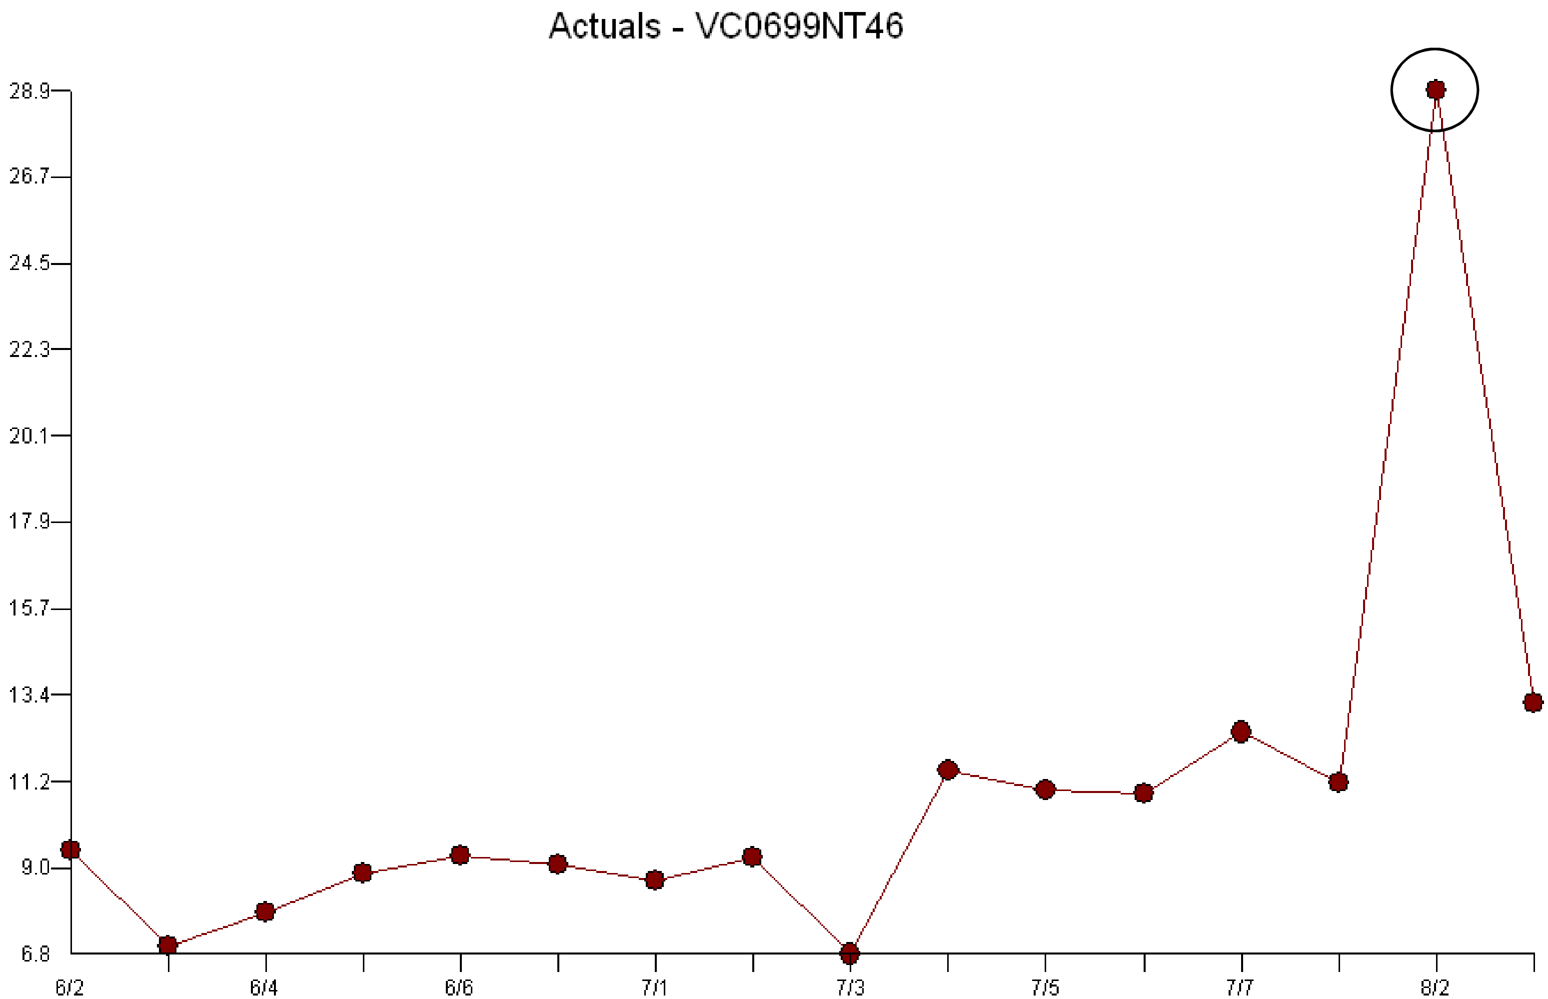

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Seasonal Pulses

- **The outlier is really a series of outliers called a seasonal pulse**

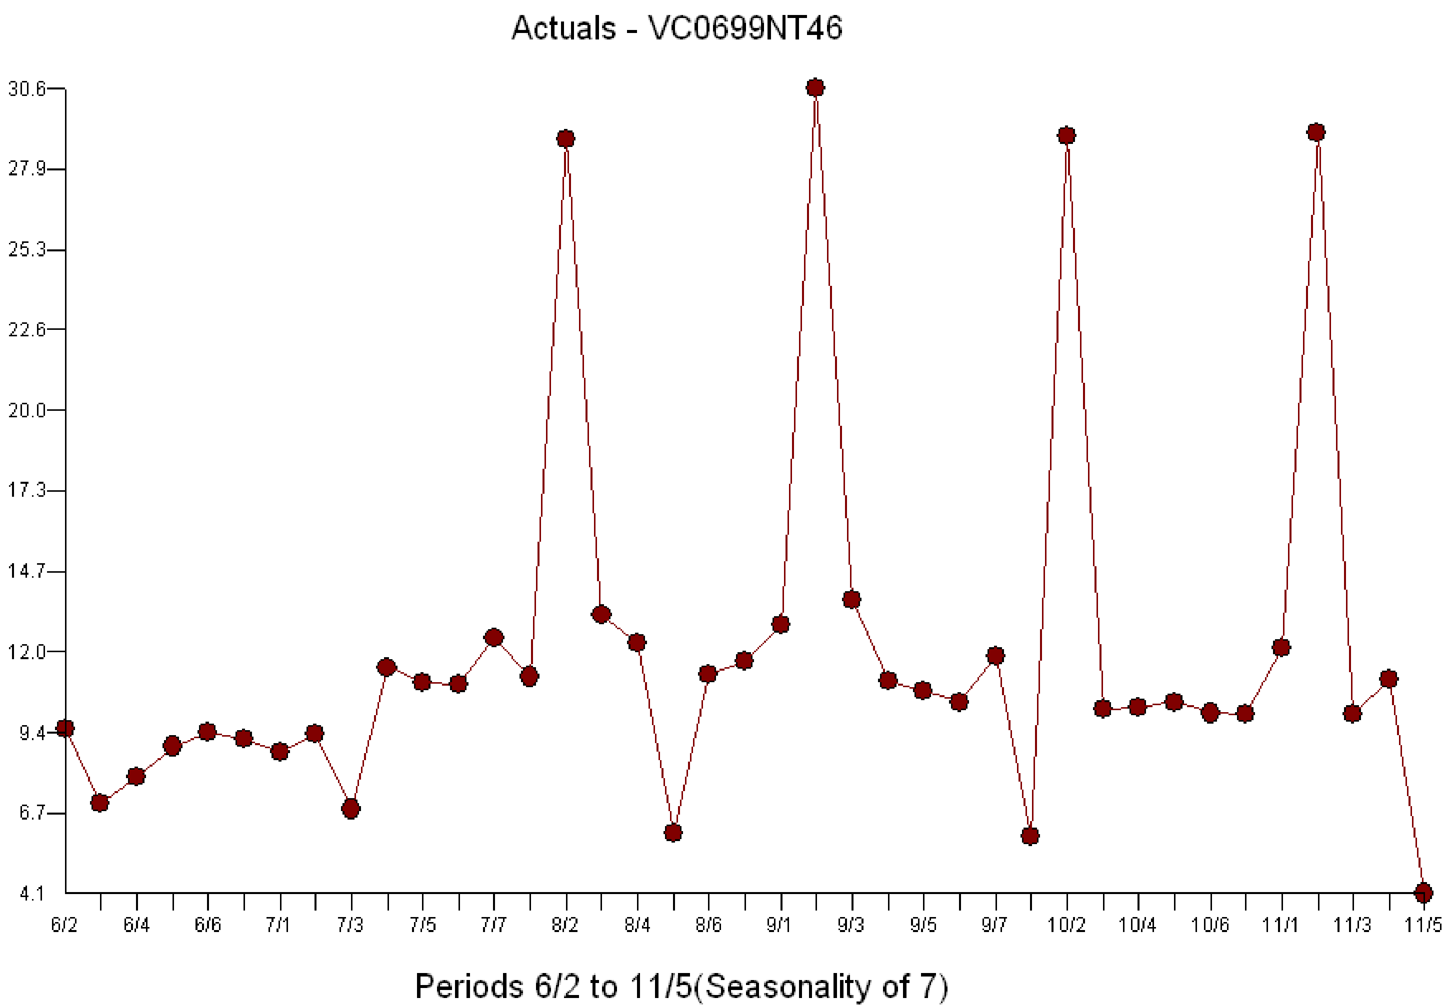

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

The "Airline Series"

- One of the most studied time series is the International Airline Passenger's series(in thousands) for monthly data from 1949 through 1960.

- Box and Jenkins didn't have the ability to detect outliers and used a log transformation to adjust the data as it seemingly had non-constant variance.

- The forecast was too high and the Box-Jenkins methodology was seen as too complicated.

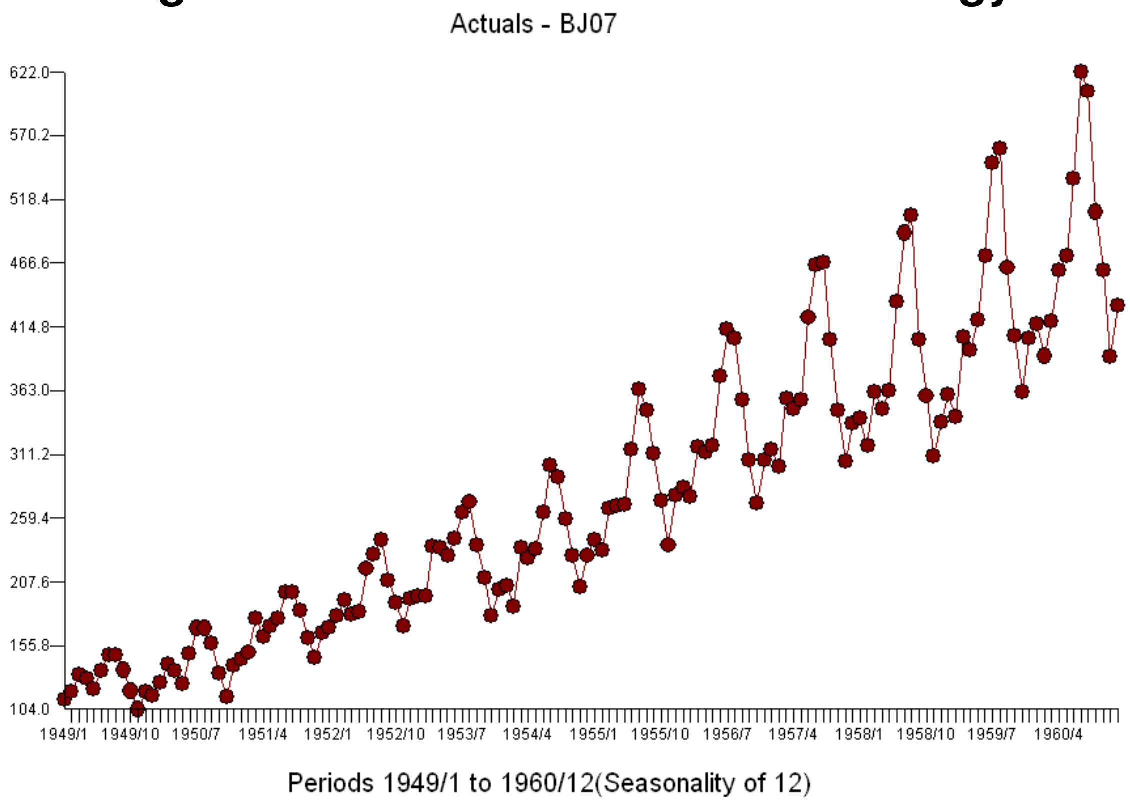

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

The "Airline Series"

- The 144 monthly observations were broken into 12 buckets (years) and they calculated the local means (assuming a model) and standard deviations for each bucket (year).

- The conclusion was that the standard deviation was increasing with the mean when it was really outliers in the last year that were skewing the situation by enlarging the standard deviation.

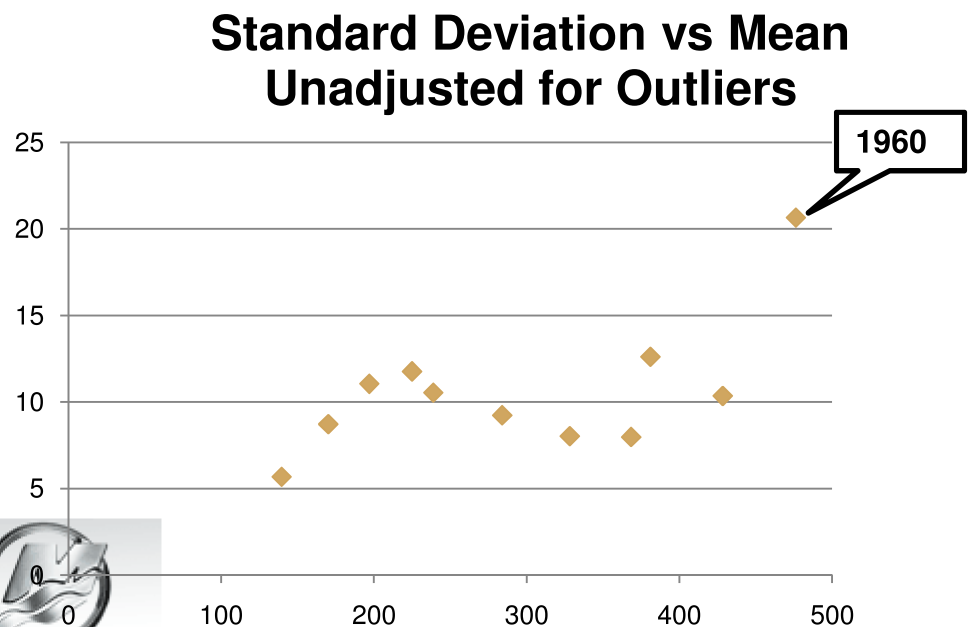

An example of "spurious correlation"

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

The "Airline Series"

- If we then fit the "airline model"(seasonal differencing and an AR1), identifying and including five outliers (three of them in the last year) we can then use the residuals to calculate the standard deviation for each of the buckets. We then plot the standard deviations against the local means of the observed series and we get another story altogether.

- The conclusion was that the standard deviation was increasing with the mean when it was really outliers in the last year that were skewing the situation.

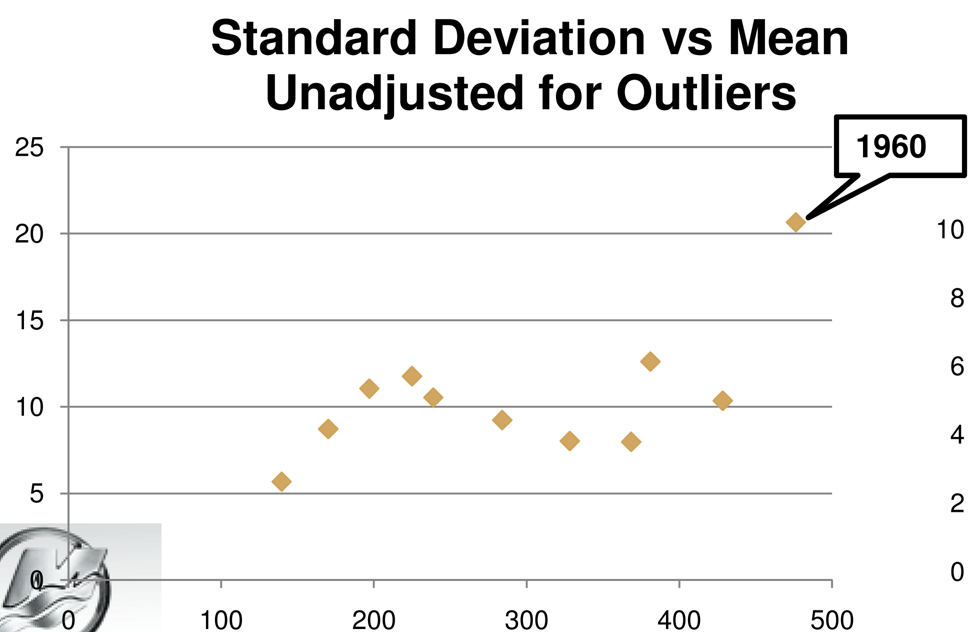

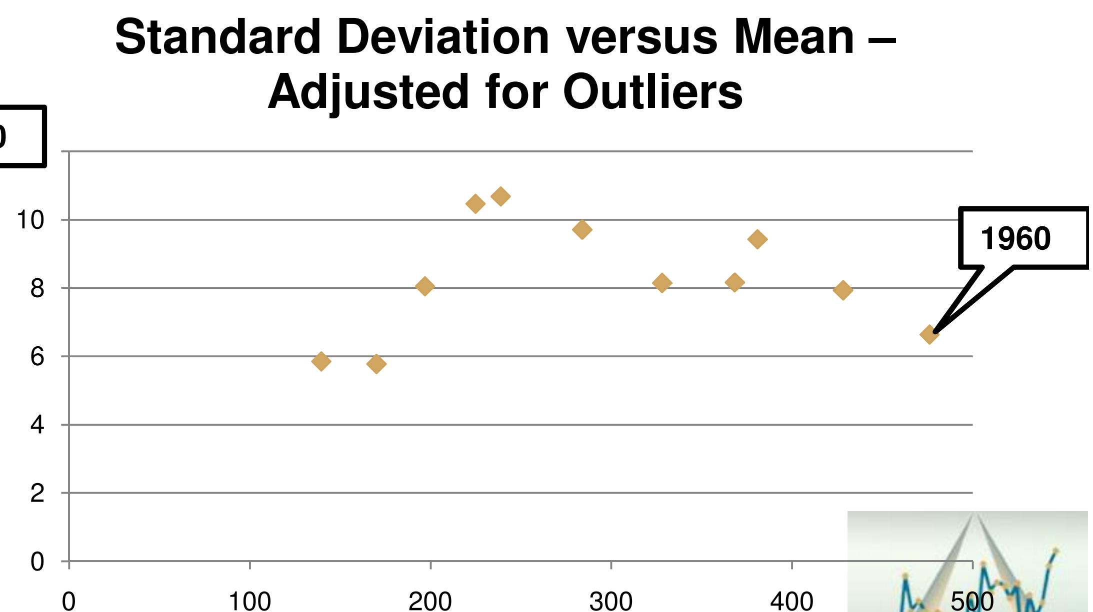

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Did you spot the outliers in 1960?

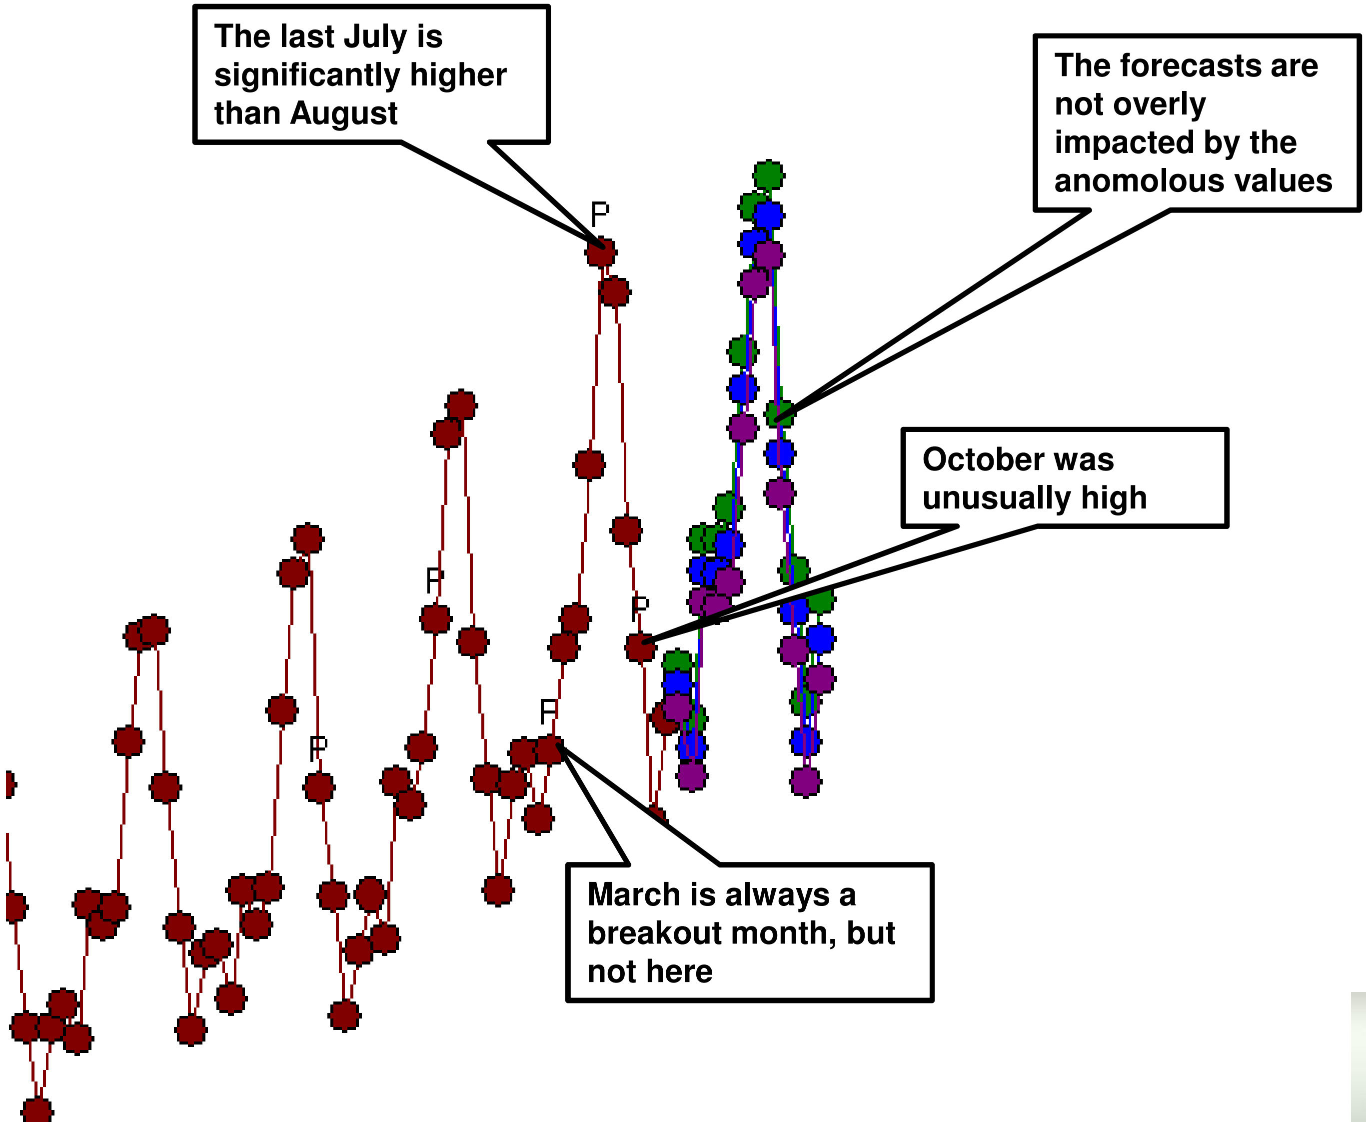

Bad Forecasting Practices still to be found in 2009

- **There are software firms that don't know that instead of taking logs as Box-Jenkins recommended in their 1976 text book, is that log transformations may not be necessary once the data has been cleansed. (Note: Some of today's textbooks also do not know or practice intervention detection !)**

- **SAS**

  - **http://support.sas.com/documentation/cdl/en/etsug/60372/HTML/default/etsug_arima_sect056.htm**

Example 7.2 Seasonal Model for the Airline Series

The airline passenger data, given as Series G in Box and Jenkins (1976), have been used in time series analysis literature as an example of a nonstationary seasonal time series. This example uses PROC ARIMA to fit the airline model, ARIMA(0,1,1)×(0,1,1)$_{12}$, to Box and Jenkins' Series G. The following statements read the data and log-transform the series:

```
title1 'International Airline Passengers';
title2 '(Box and Jenkins Series-G)';
data seriesg;
   input x @@;
   xlog = log( x );
   date = intnx( 'month', '31dec1948'd, _n_ );
   format date monyy.;
datalines;
```

Bad Forecasting Practices still to be found in 2009

- **There are software firms that don't know that instead of taking logs as Box-Jenkins recommended in their 1976 text book, is that log transformations may not be necessary once the data has been cleansed. (Note: Some of today's textbooks also do not know or practice intervention detection !)**

- **Oracle**
  - **http://oracledmt.blogspot.com/2006/03/time-series-forecasting-2-single-step.html**

First we stabilize the variance. This can be done by applying a Box-Cox power transform. This transform has the following form: $y(h) = (y^h - 1) / h$, if h is not equal to 0 and $y(h) = \log(y)$ if h is 0. In general, the LOG transform (h=0) is a good choice for removing increasing variability. Figure 2 shows the transformed series after the LOG transform. The upward trend over time is still visible but the amount of variation in the series is about the same throughout the series.

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Bad Forecasting Practices still to be found in 2009

- There are software firms that don’t know that instead of taking logs as Box-Jenkins recommended in their 1976 text book, is that log transformations may not be necessary once the data has been cleansed. (Note: Some of today’s textbooks also do not know or practice intervention detection!)
- XLSTAT
    - http://www.xlstat.com/en/support/tutorials/arima.htm

We notice on the chart, that there is a global upward trend, that every year a similar cycle starts, and that the variability within a year seems to increase over time. Before we fit the ARIMA model, we need to stabilize the variability. To do that, we transform the series using a log transformation. We can see on the chart below that the variability is reduced.

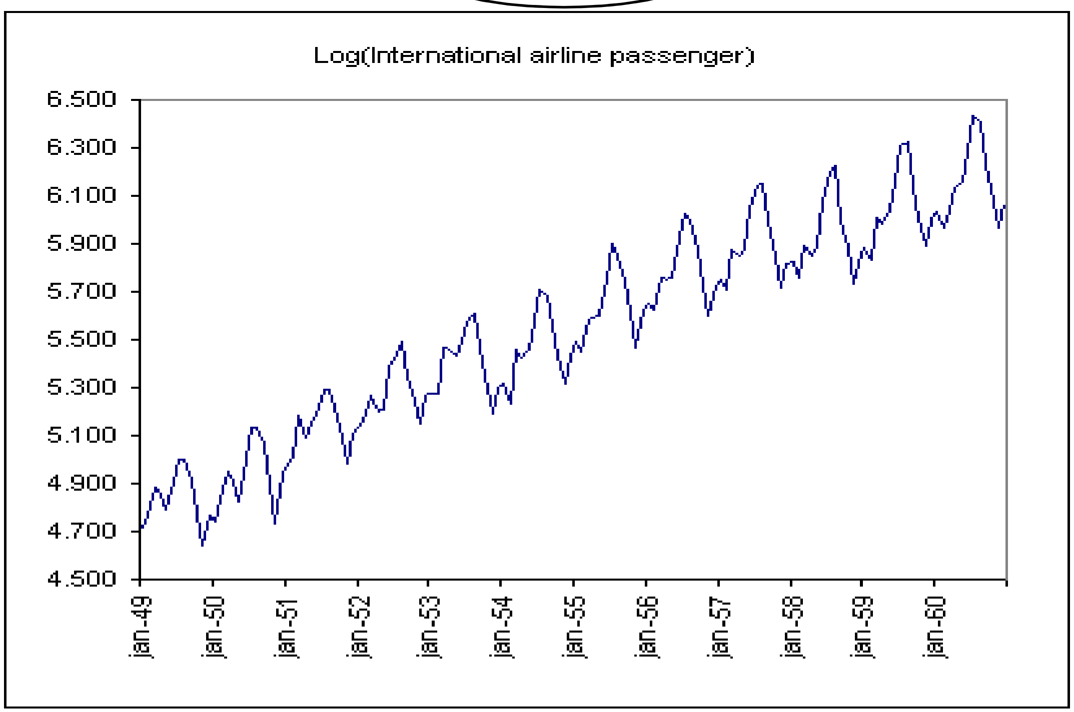

Bad Forecasting Practices still to be found in 2009

- There are software firms that don't know that instead of taking logs as Box-Jenkins recommended in their 1976 text book, is that log transformations may not be necessary once the data has been cleansed. (Note: Some of today's textbooks also do not know or practice intervention detection !)
- Visual Numerics (IMSL- International Math and Statistics Library)
  - http://www.vni.com/products/imsl/documentation/fort06/stat/NetHelp/default.htm?turl=bctr.htm

Example 1

Consider the Airline Data (Box and Jenkins 1976, page 531) consisting of the monthly total number of international airline passengers from January 1949 through December 1960. Routine BCTR is used to compute a forward Box-Cox transformation of the first 12 observations. In the transformation SHIFT and POWER are each set to zero, which corresponds to taking natural logarithms of the data.

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Bad Forecasting Practices still to be found in 2009

- There are software firms that don't know that instead of taking logs as Box-Jenkins recommended in their 1976 text book, is that log transformations may not be necessary once the data has been cleansed. (Note: Some of today's textbooks also do not know or practice intervention detection!)

- MATLAB

  - http://www.mathworks.com/products/statistics/demos.html?file=/products/demos/shipping/stats/stattsdemo.html

Examine Trend and Seasonality

This series seems to have a strong seasonal component, with a trend that may be linear or quadratic. Furthermore, the magnitude of the seasonal variation increases as the general level increases. Perhaps a log transformation would make the seasonal variation be more constant. First we'll change the axis scale.

```
set(gca,'YScale','log');
```

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Bad Forecasting Practices still to be found in 2009

- There are software firms that don't know that instead of taking logs as Box-Jenkins recommended in their 1976 text book, is that log transformations may not be necessary once the data has been cleansed. (Note: Some of today's textbooks also do not know or practice intervention detection!)
- Mathematica
  - http://media.wolfram.com/documents/TimeSeriesDocumentation.pdf

We see that the variance, riding on the trend, is also changing with time. We need to transform the series into a constant variance before modeling it further. To stabilize the variance, a nonlinear transformation such as a logarithmic or square-root transformation is often performed. In this example, we try a natural logarithmic transformation, $y_t = \ln x_t$.

This is the time plot of the airline data after the logarithmic transformation. Note that Log[*aldata*] gives the logarithm of each of the entries of aldata.

```
In[40]:= ListLinePlot[Log[aldata], AxesLabel -> {"t", "ln(xt)"}]
```


MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Bad Forecasting Practices still to be found in 2009

- There are software firms that don't know that instead of taking logs as Box-Jenkins recommended in their 1976 text book, is that log transformations may not be necessary once the data has been cleansed. (Note: Some of today's textbooks also do not know or practice intervention detection!)

- Stata

  - http://www.stata.com/bookstore/pdf/arima.pdf

▷ Example 3: Multiplicative SARIMA model

One of the most common multiplicative SARIMA specifications is the $(0,1,1) \times (0,1,1)_{12}$ "airline" model of Box, Jenkens, and Reinsel (1994, sec. 9.2). The dataset `airline.dta` contains monthly international airline passenger data from January 1949 through December 1960. After first- and seasonally differencing the data, we do not suspect the presence of a trend component, so we use the `noconstant` option with `arima`:

```
. use http://www.stata-press.com/data/r10/air2
(TIMESLAB: Airline passengers)
. generate lnair = ln(air)
. arima lnair, arima(0,1,1) sarima(0,1,1,12) noconstant
```


MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Outlier Detection – Pulse

- **Pulse – Fire in the warehouse in April (0,0,0,0,0,0,0,0,1,0,0,0,0,0)**

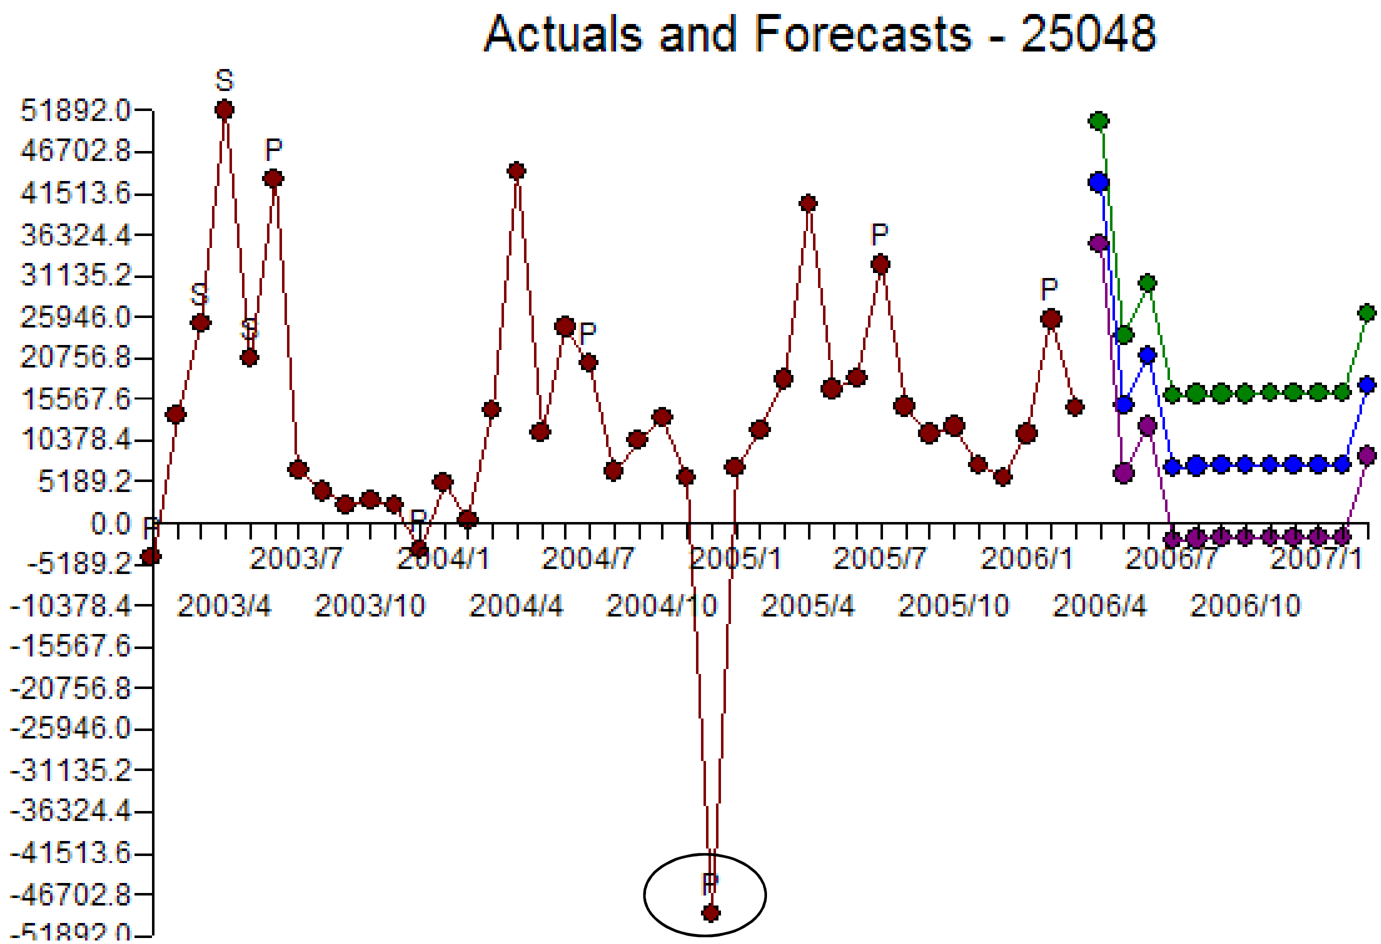

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Outlier Detection – Seasonal Pulse

- **Seasonal Pulse – February emerges later during the year (0,1,0,0,0,0,0,0,0,0,0,0,0,1)**

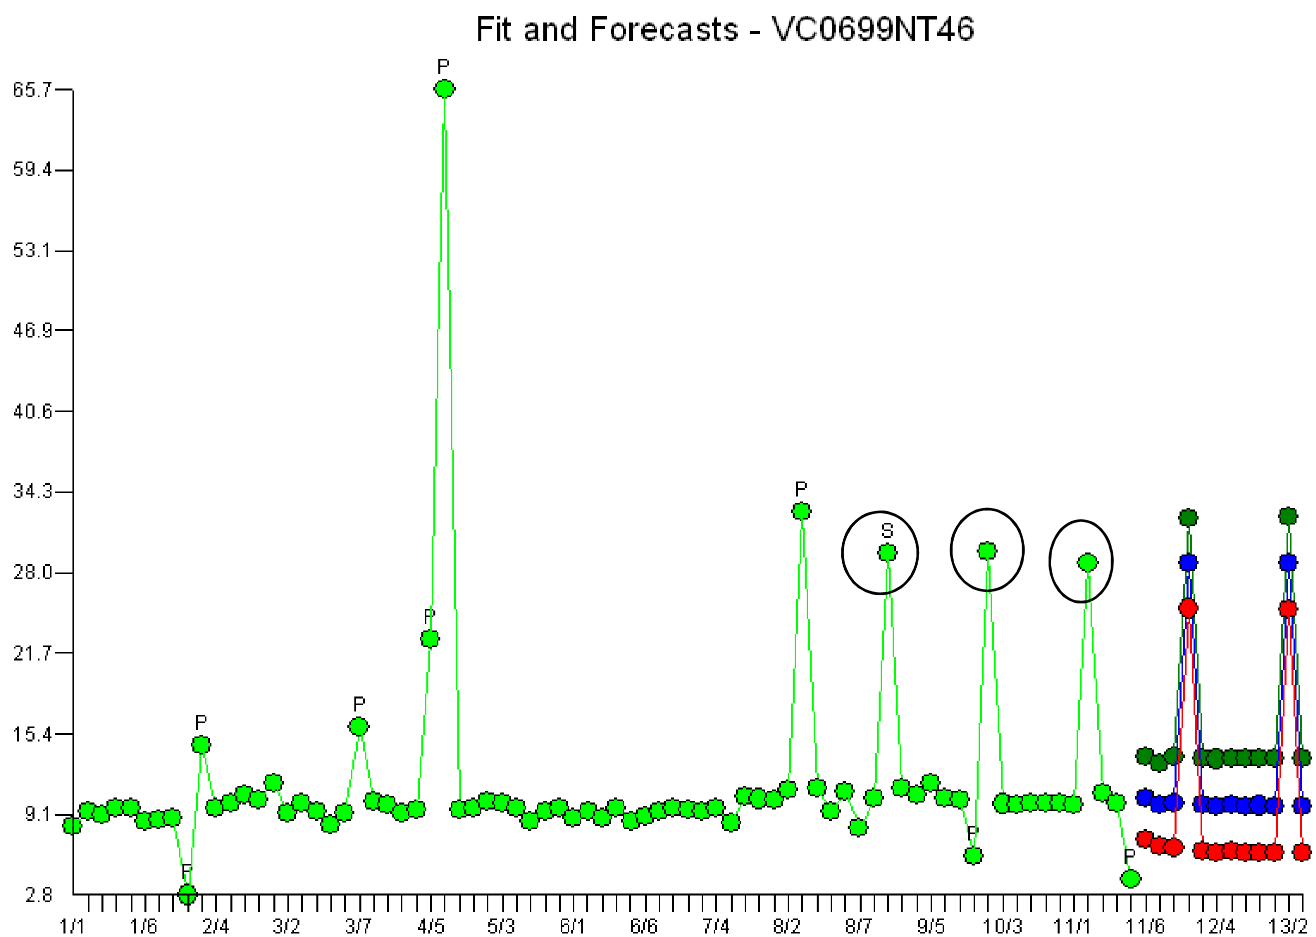

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Tough Series to Model

- From a visual it looks like a seasonal model that is increasing, right?

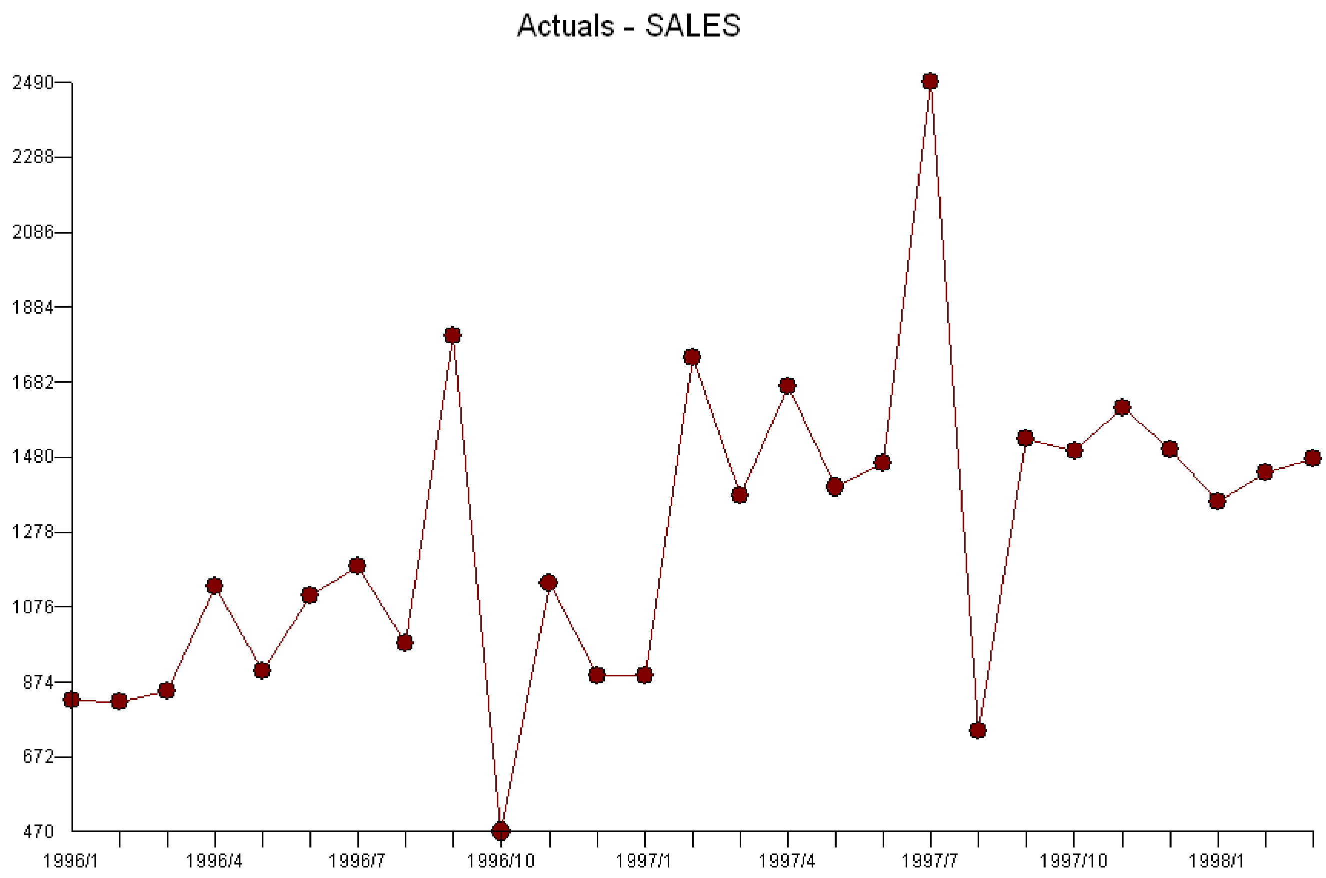

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

What Would PROC Do?

- Drive a Holt-Winters Model through its heart and predict an upwards trend?
- Be fooled that there is seasonality even though the blips are 13 periods apart not 12?

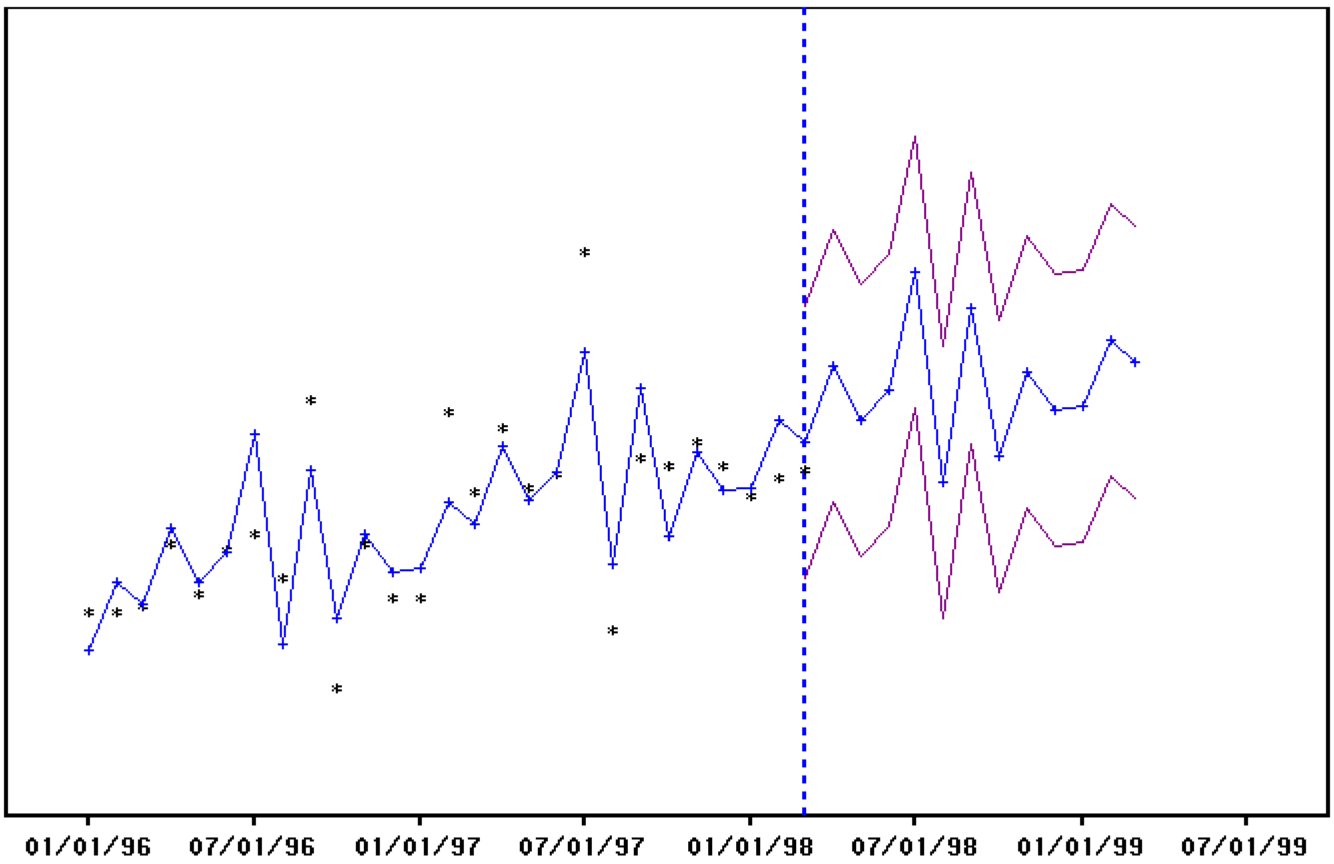

Outlier Detection – Level Shift

- **Level Shift –Competitor drops out of the market and an 'one-time' increase in market share gain (0,0,0,0,1,1,1,1,1,1,1,1)**

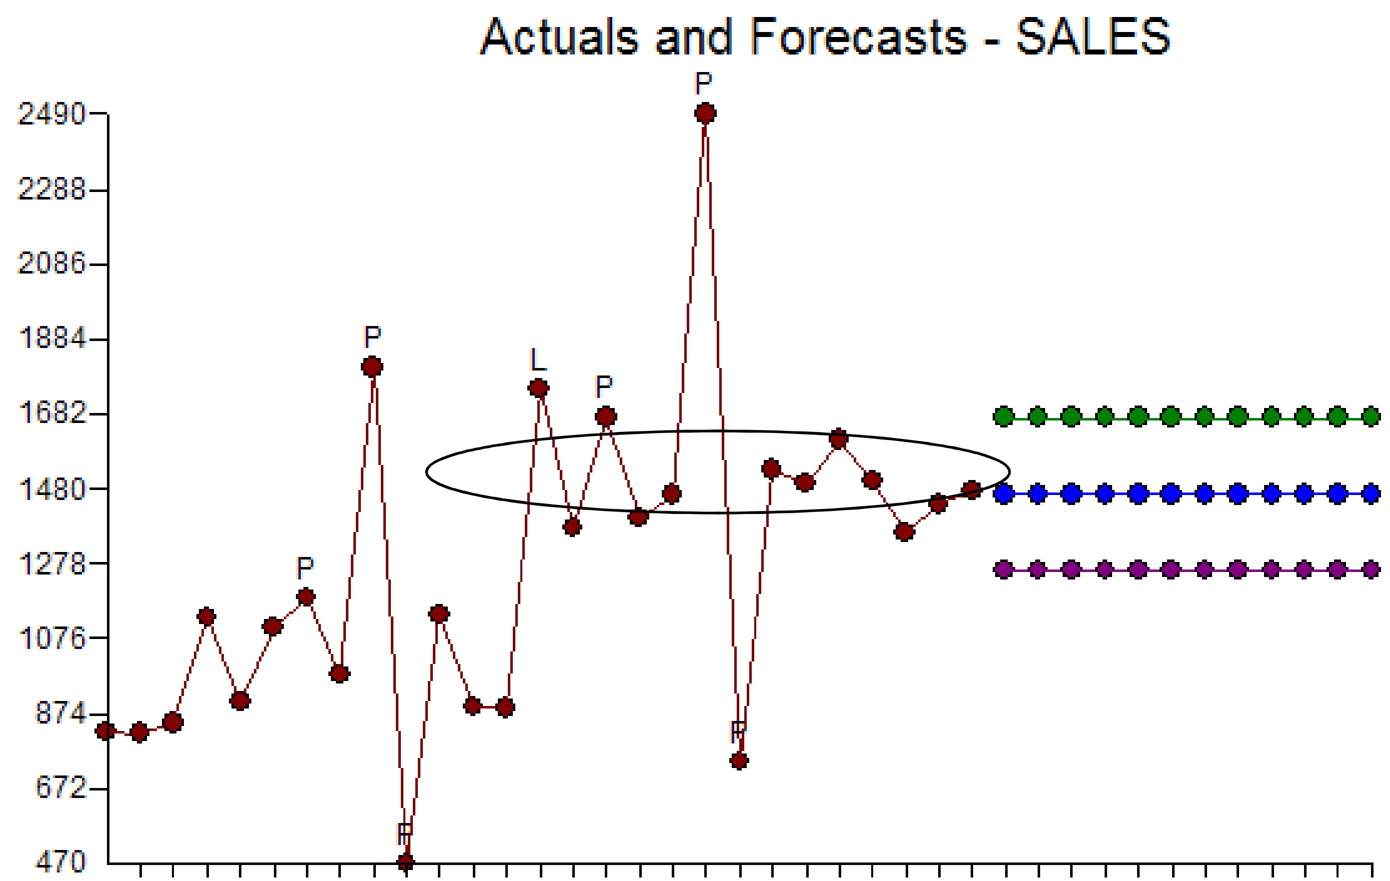

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Outlier Detection – Local Time Trend

- **Local Time Trend – A new trend up or down very different from the past (0,0,0,0,1,2,3,4,5,6,7,8,9,etc.)**

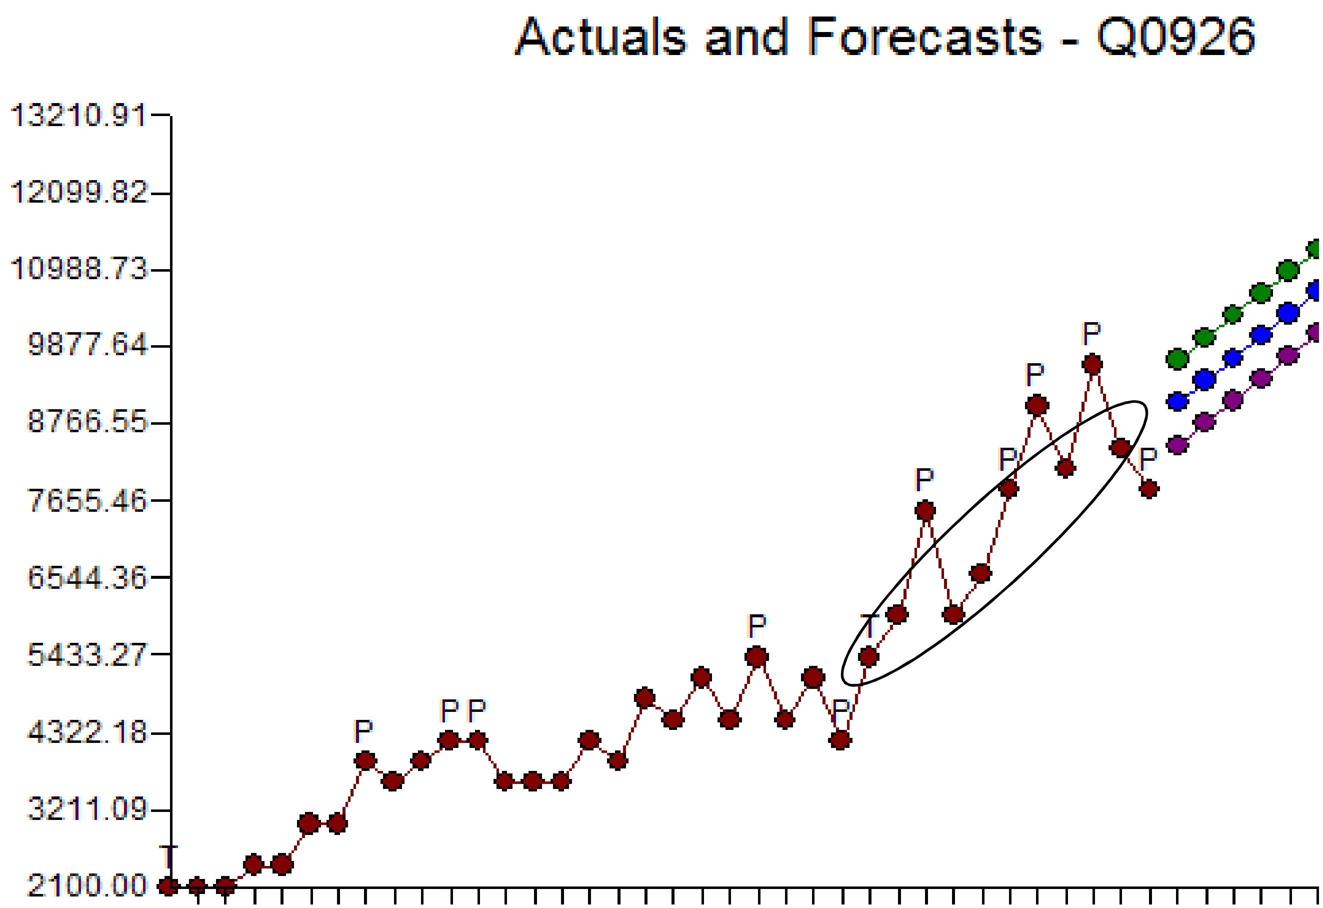

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Outlier Detection – What should you do about it?

- **User Provides knowledge - before the modelling process begins – If there is some domain knowledge that there was an event in the past then this information should be included in the model as a possible input variable so that the observed value is not "adjusted". In this case an actual variable now has a coefficient and can explain the impact so that it's effect can be anticipated in the future if the candidate variable is operational or in effect.**

- **Action - You don't want to believe a pulse and you should adjust the pulse to "where it should have been" thus providing a robust estimation of the model parameters.**

- **No Action - If you do not adjust for outliers then the coefficients in the model will be skewed creating a false image of the systematic behavior. The forecast may be higher or lower than anticipated. The causal relationship may have an incorrect snapshot of the relationship between price and sales for example.**

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Outlier Detection - Out of Model – The Downside

- **<u>Keeping</u> what is unusual within the standard deviation bound:**
  - **If there is a level shift up in the last 5 periods, it will skew the standard deviation upwards and therefore the values will not be considered unusual**

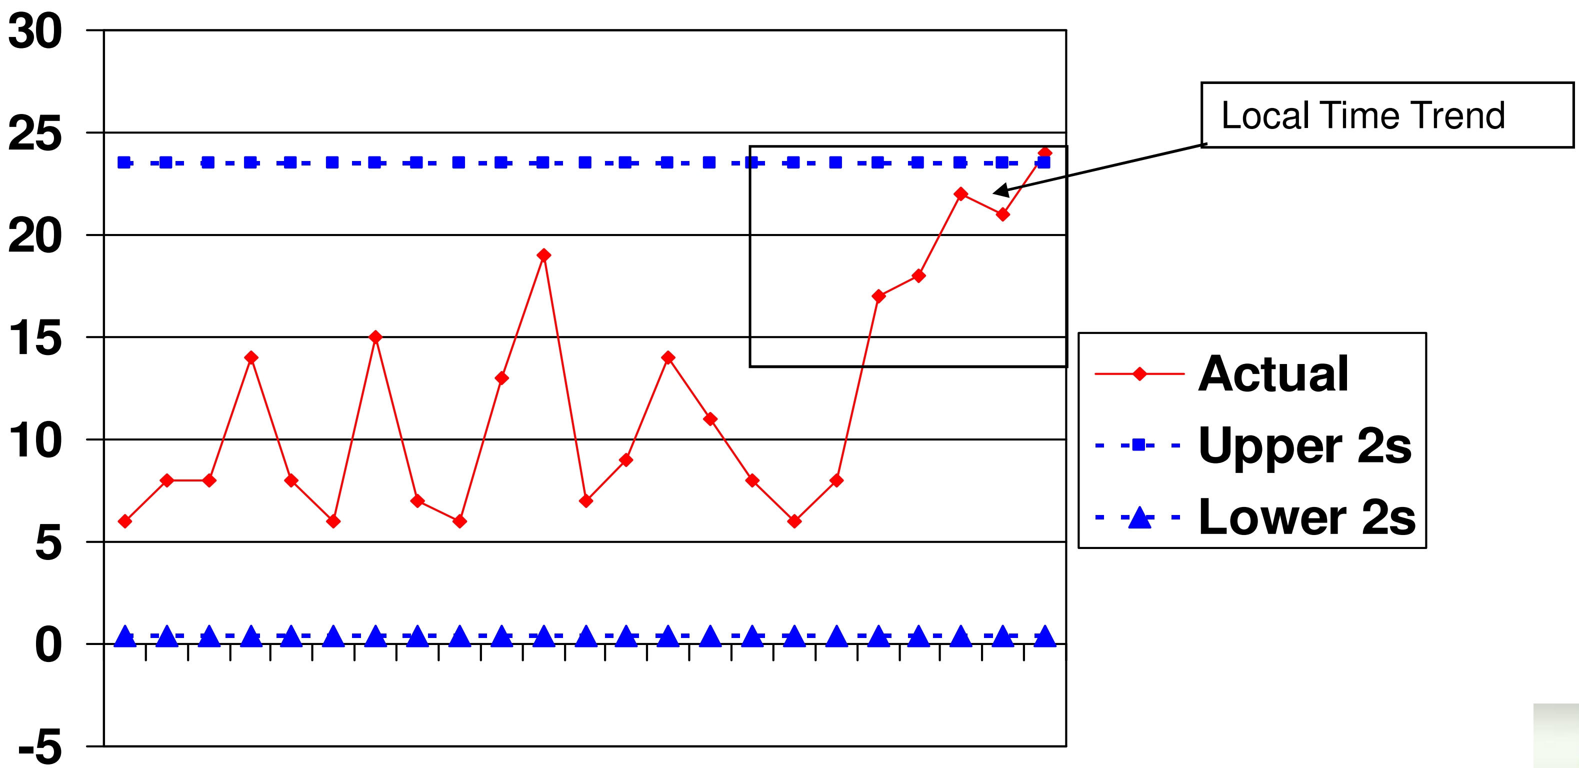

Outlier Detection - Out of Model – The Downside

- **<u>Keeping</u> what is unusual within the standard deviation bound:**
  - **If a pulse occurs that is an "inlier" and not outside the standard deviation bounds then it will not be identified as you need a model first to identify this situation (e.g. 1,9,1,9,1,9,5)**

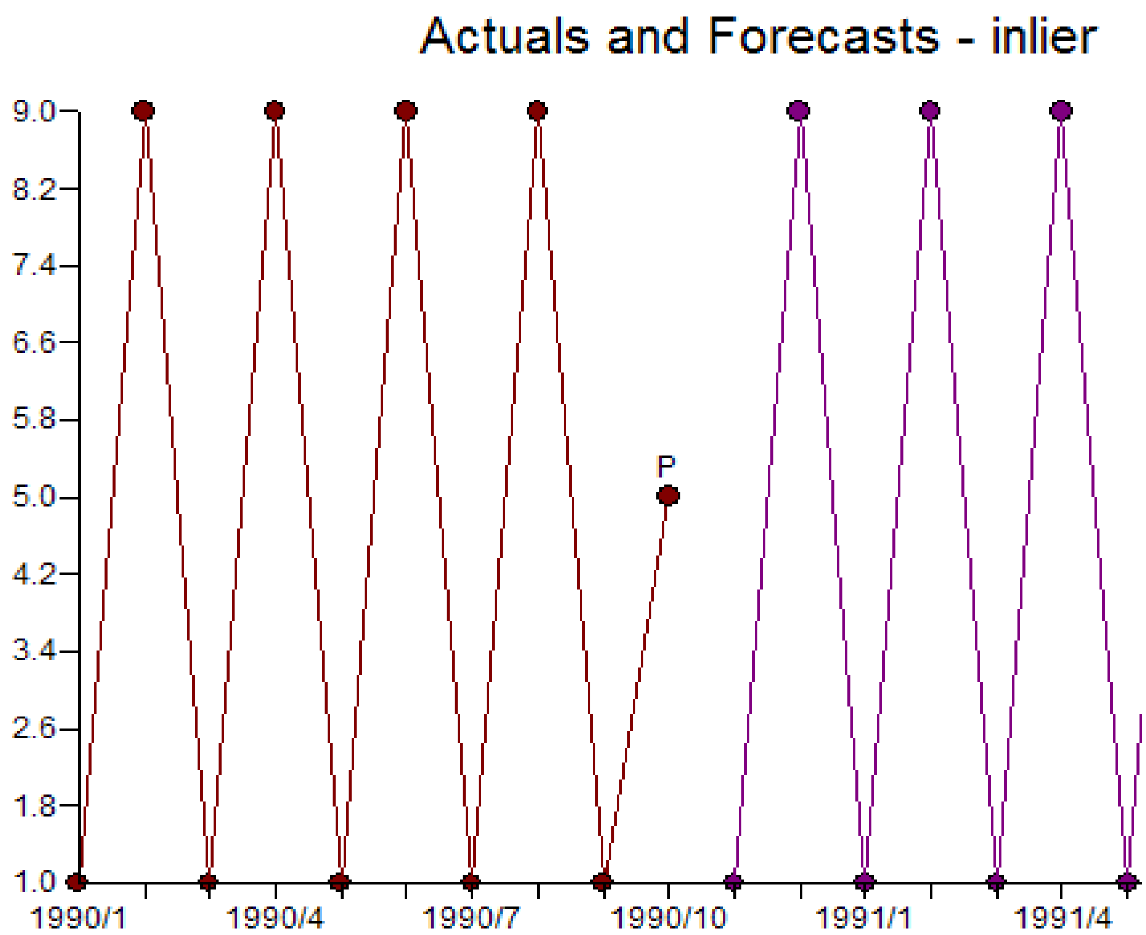

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Outlier Detection - Out of Model Approach

- Calculate Residuals (Actual - Mean)
- Calculate the standard deviation around the mean
- Specify the # of standard deviations around the mean that will be considered an outlier (e.g. 3)
- Identify those observations outside the standard deviation
- Replace the unusual observations with the mean or interpolation
- Specify the number of iterations to go through the outlier removal process

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Outlier Detection - Out of Model – The Downside

- **The mean is skewed by the outliers and when the standard deviation is calculated it is larger than what it should be which will cause more observations to be removed than necessary**

- **Series with positive correlation have a understated standard deviation and while the negatively correlated series have an overstated standard deviation.**

- **Everything is an outlier! – Some software allows you to recursively adjust the outlier and if you do then you may be adjusting outliers which were really not outliers.**

- **If a level shift occurs and is not accounted for then the forecast will be "off" due to the previous historical data affecting the forecast. The level shift is not outside the standard deviation bound, but certainly is a big change in the process.**

- **If the outlier detection scheme is looking only for one type of outlier and may <u>remove</u> data that was in fact real:**

    - **Seasonal Pulses may be high/low and very real as they occur regularly or worse using a seasonal model (AR12) when it is in fact not.**

Outlier Detection - Within Model Approach

- Identify a possible model for the data.

- Identify the outlier in the presence of the model effect.

- A statistically valid test is performed on the unusual value which uses a “standard deviation of the error” that excludes the impact of the outlier and therefore is more robust

- Consider identifying outliers first and then the model form.

- Evaluate the alternatives of model first then outlier detection. Compare results and determine the optimal strategy to follow for each dataset.

A Critique of Automatic Forecasting Software

REPORT CARD

________	C+
________	B-
________	A

Forecasting Methods Family Tree

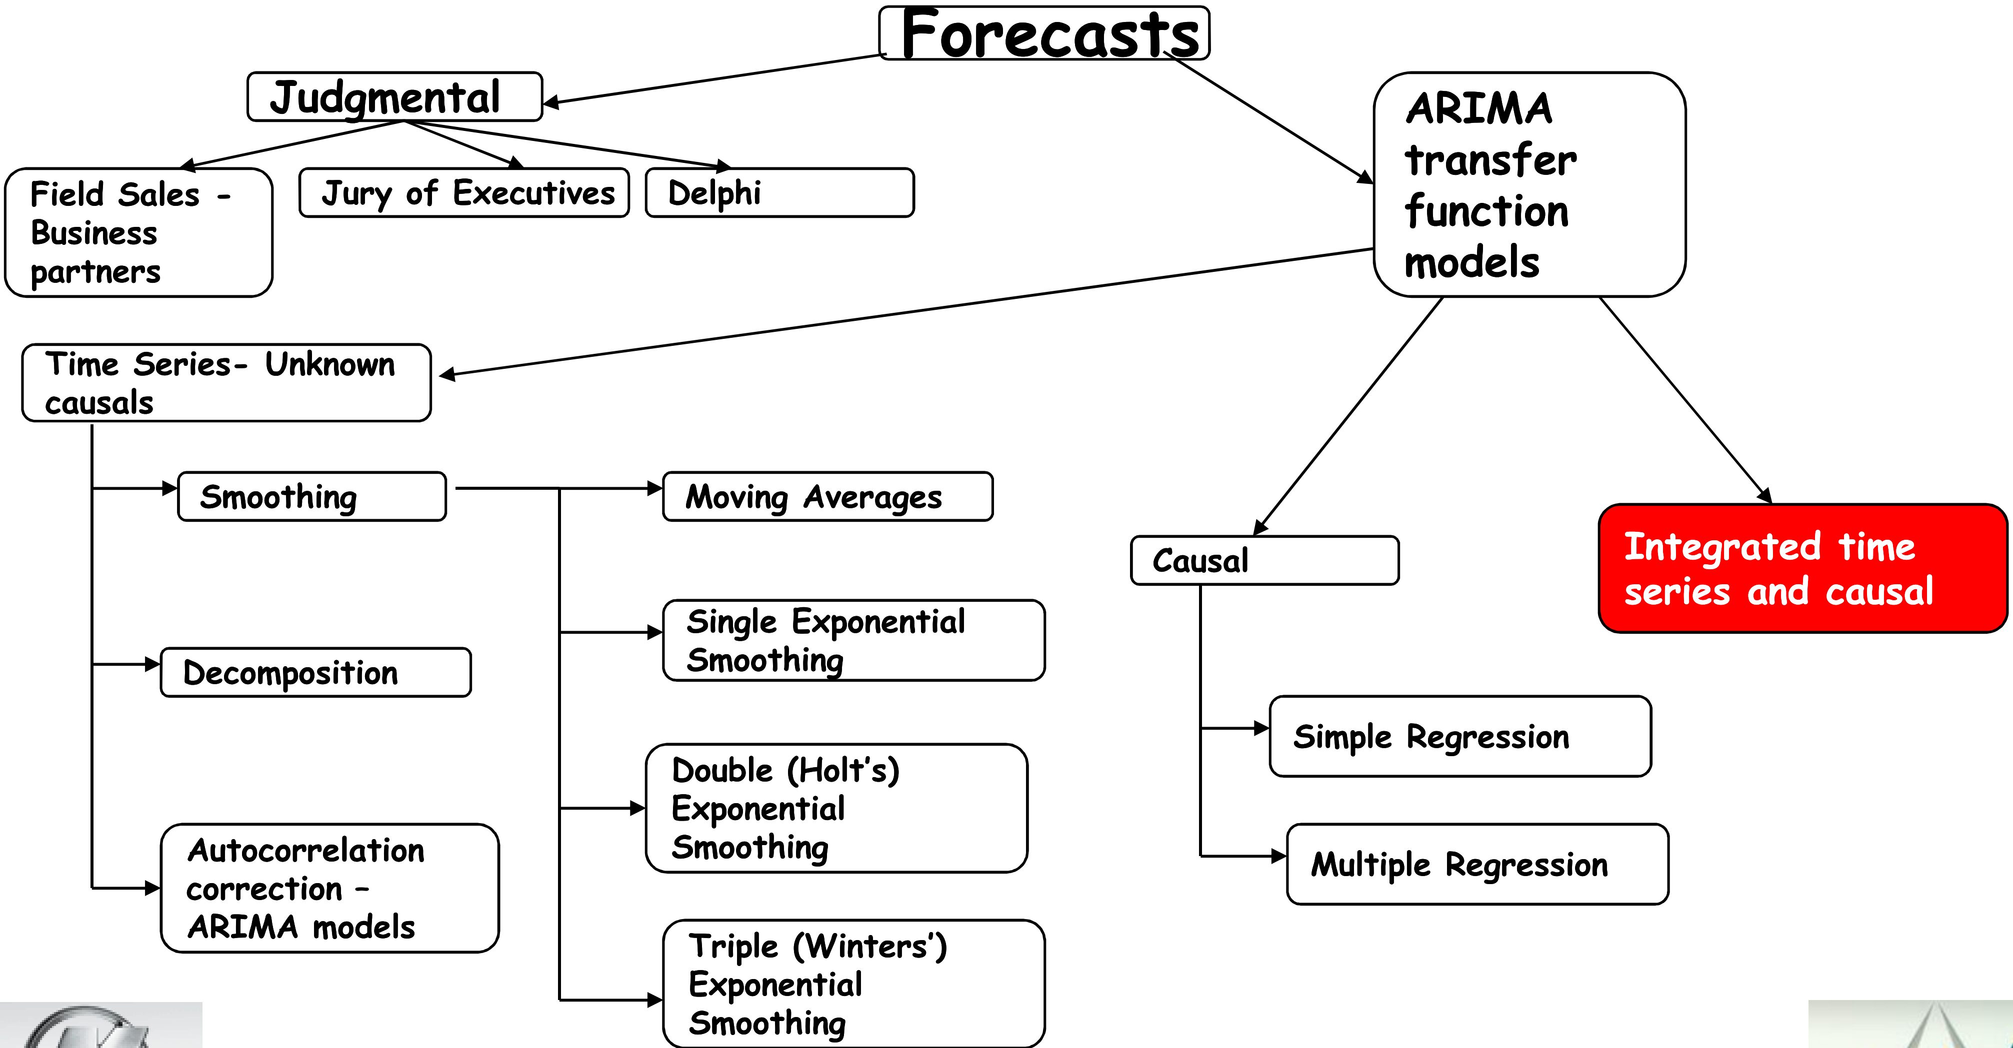

MERCURY MARINE

38
AUTOMATIC FORECASTING SYSTEMS

History of Modelling

- 1795 - Legendre – Regression - Developed for cross-sectional data and later abused using with time series data – You can swap first observation with last observation and get the same answer
- 1920 - Seasonal Decomposition
- 1927 - Slutsky-Yule – Identified that applying a moving average to a random process may in itself create a pattern when no existed previously
- 1944 - Brown – Exponential Smoothing
- 1957 - Holt – Holt Method
- 1960 - Holt/Winters – Holt/Winters
- 1960 - Chow test – Parameter Changes
- 1965 - Almon – Polynomial Distributed Lag
- 1967 - Hiskin/Young, Musgrave – Seasonal Adjustment Census Bureau
- 1970 - Box & Jenkins – Time Series Analysis Textbook –
  - Introduced a Generic model form which *__all__* models are a subset
  - Introduced a "Data based" approach of building a unique model and coefficients for each data set
- 1976 - Box and Tiao – Interventions
- 1988 - Chang, Tiao and Chen – Innovational Outliers
- 1988 - Tsay – Level Shifts, Variance Change

"Model based" assuming the relationship within the data is a certain weighting scheme and the number of periods to weight.

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Progress in a Real World View of History of Modeling

- Let's use the last 100 days to predict tomorrow's rainfall using an average
- Let's only use the last 12 days and weight the more recent data more and the older data less
- Let's use only the last 7 days using a weighting scheme and provide a bump up on Friday's as it rains more on Friday's
- Let's use an "additive" method to adjust by way off addition or subtraction for the forecast fluctuations
- Let's use a "multiplicative" method to adjust by way off %'s up and down for the forecast fluctuations
- By using the Box-Jenkins approach of calculating lags of the history using regression to identify the length of time to use and the weighting instead of assuming
- By using Intervention Detection, you can add unspecified causal variables to the model that adjust for outliers, level shifts, local time trends and seasonal pulses that if not accounted for will distort the coefficients in the model and thereby the forecasts

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Box-Jenkins Modeling Steps

Identification - Calculate statistics on the data to suggest a model form (length and weight) using the ACF & PACF

Estimation - Taking suggested model form and estimating the optimal coefficients

Diagnostic Checking - Making sure that the residuals are constant mean/variance, random and no autocorrelation

Forecasting - Take the estimated model and generate X period out forecasts

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Regression Modelling

- **Causal variables - Variables like Price might have a lead or lag relationship and that exact period may be difficult to identify. Assuming it is not going to help you.**

- **Dummy variables – Outliers need to be provided and/or identified and adjusted for by the system.**

- **Memory – There is a period to period relationship that exists in the data. The historical data implicitly captures the effect of omitted causal variables.**

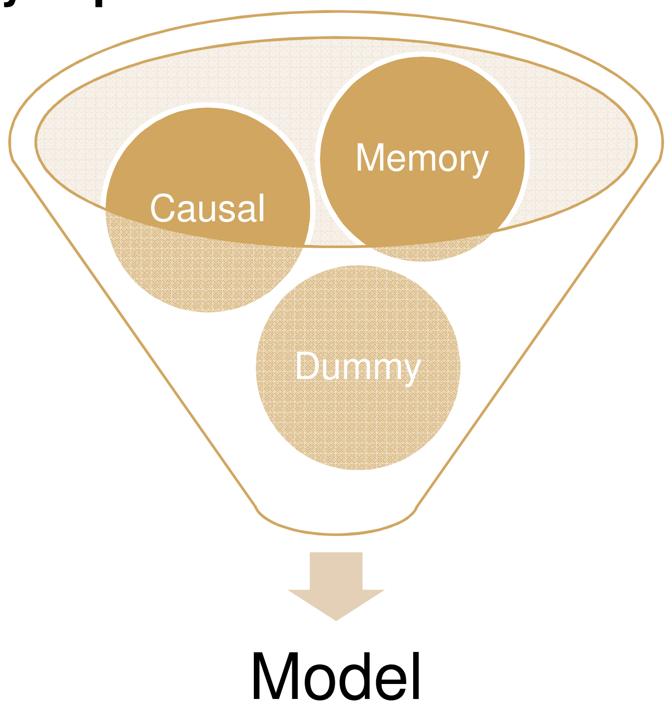

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Regression Modelling

- Most Software packages allow a user to provide causal variables like price and promotion and events (e.g. holidays, outliers), but assuming that the relationship between the causals and the output series are contemporaneous

- "Skipping Identification" and going right into Estimation means that the lag or lead relationship between the causals and the output series has not been attempted
  - Some will assume a theoretical 3 month lag relationship, but is it?
  - Some will attempt to plot lags of causals vs. the output variable to "see" the lag or lead relationship
  - Some with more knowledge will use methods to use a statistic called the "cross-correlation function" to identify the relationship
  - Some may review the residuals from the model and then add in an ARIMA of lag 1 to correct for period to period relationship of the output variable (e.g. Hildreth-Liu or Cochrane-Orcutt) instead of identifying the model

- Trying to Identify the relationship takes more computing time
  - Is there a lead or lag relationship?
  - Are there are outliers that need to be incorporated?

Automatic Modelling

- User might be asked(forced?) to order the data into groupings before any modelling ever occurs so the modeling process doesn't get "fooled":
  - Is the data seasonal?
  - Is there intermittent demand?
  - Do you have any events (promotions, interventions)?

- System tries various quadratic equations to get the best fit, but lacks any ability to forecast.

- User specifies the model and the system estimates the optimal coefficients.

- User specifies which of the different criteria(i.e. smallest AIC, BIC, SIC, RMSE etc.) to be used to determine the "goodness of fit" from a pre-specified list of models using a withheld number of observations yielding strange results like seasonal models when there is no seasonality in the data at all.

- <u>Heuristics</u> determine model form, variables that are significant and suggests and includes interventions into the model.

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

MODELLING VS FITTING

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Model Fitting

- Fast and easy to do, but not likely to match the “fingerprint to the killer”
- Fitting a round peg into a square hole 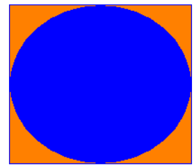
- Can a “pick-best” approach work in an infinite sample space?
  - The “fitters” will take a list of ~10, 25, 50 models and try to find the model that best “fits” the data. The process will then tell you that it is “optimizing” the parameters, but it just can’t be as you assumed a model to begin the process
- Sometimes the “fitters” model get fooled by only partially describing the data
  - Seasonal dummy model used with data that only has <u>some</u> seasonal months
  - A level shift is thought to be a series with an upward trend

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

"Gaming the System" by Way of Withholding Data

- **We could find which model and which number of observations that are withheld would result in the best fit if we really wanted to, but should you?**
- **If you have ~50 models and fit them using 1 period withheld**
- **If you have ~50 models and fit them using 2 periods withheld**
- **And so on until you have done up until 12 periods withheld**
- **You would have the combination of one type of model and a certain number to be withheld that would be the "winner" of the best outcome (e.g. model 17 and 4 withheld)**
- **The reality is that while this exercise in futility would certainly yield the smallest fitting statistic, but it really is not capturing the pattern in the data and is only an exercise "mathematical manipulation"**

More on Withholding Observations

- Some will allow users to specify the number of observations to withhold to allow the model “train” so that it optimizes the model for this withhold set of data
- It is the case of the tail wagging the dog

- How do you know how many observations to withhold and what happens if I change the withhold from 6 to 7, will my model and forecast change? You betcha!
- This approach builds a model on data and then changes the coefficients based on the most recent data.
- What if there are outliers not captured by the process in the withhold data? It will skew the model and forecast
- Are the older data worthless? They are rendered so as the withheld data is used to determine the “best” coefficients

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Customized Modeling

- Much slower and requires complicated schemes to sift through the patterns in the data to build a customized model for the each data set.

- Did it get the model 100% right? Probably not, but then again the fitter didn't even try and took a passive rather than active approach. It's like passing to the other family in "Family Feud" instead of trying to answer the question yourself

- It's like getting a custom made suit that fits your dimensions

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

IGNORING THE ASSUMPTIONS AT YOUR OWN PERIL !

THE DEVIL IS IN THE DETAILS !

Assumptions

- The statistical test is what determines if the model has significant statistical value. In order to have a valid statistical test, you must have some assumptions met and if they are not then the statistics behind your model are ***null and void.***

- Constant Mean and Constant Variance in the Residuals. If the residuals are not random (actuals – fit = residuals) and have a pattern then you have not accounted for <u>all</u> of the pieces of the model that describe the pattern in the data, you have misspecified the model and your forecast will reflect that.

  - How to tell?
    - A Plot of the residuals
    - The ACF/PACF plot

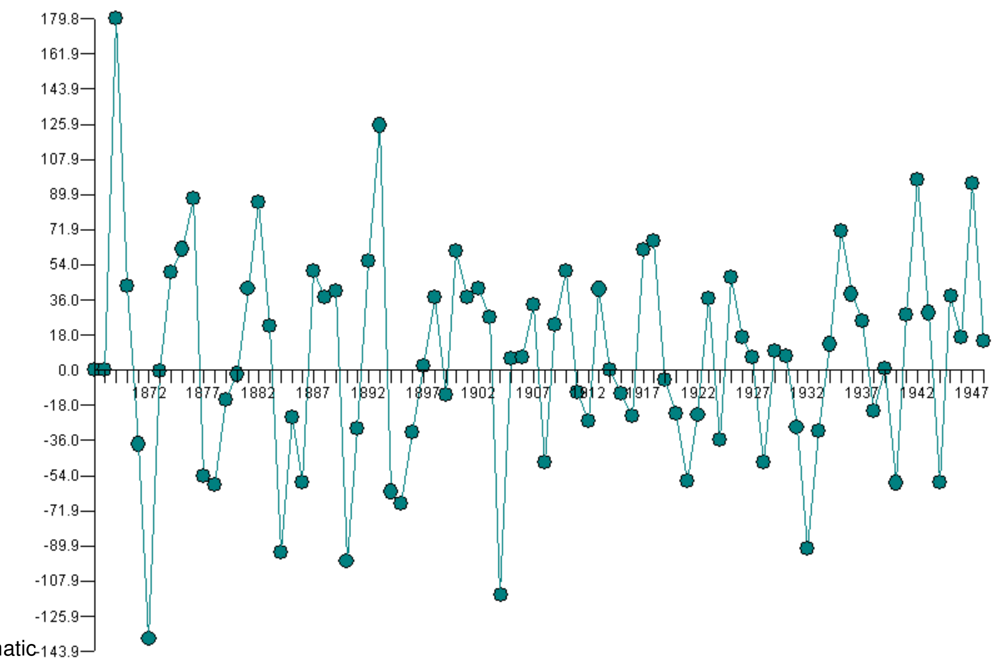

Eight Examples of Possible Violations

Mean of the Error Changes: (Taio/Box/Chang)

1. A 1 period change in Level (i.e. a Pulse)
2. A contiguous multi-period change in Level (Intercept Change)
3. Systematically with the Season (Seasonal Pulse)
4. A change in Trend (nobody but Autobox)

Variance of the Error Changes:

5. At Discrete Points in Time (Tsay Test)
6. Linked to the Expected Value (Box-Cox)
7. Can be described as an ARMA Model (Garch)
8. Due to Parameter Changes (Chow, Tong/Tar Model)

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Constant Variance Assumption

- What happens when the variance is not constant over time?

- Here is the often studied IBM Stock Price in the Box-Jenkins text book. Note that the variance increases as the level of the series decreases. A knee jerk reaction by an economists/statisticians is to "take the log of the series" which can have bad effects and can be remedied in another way.

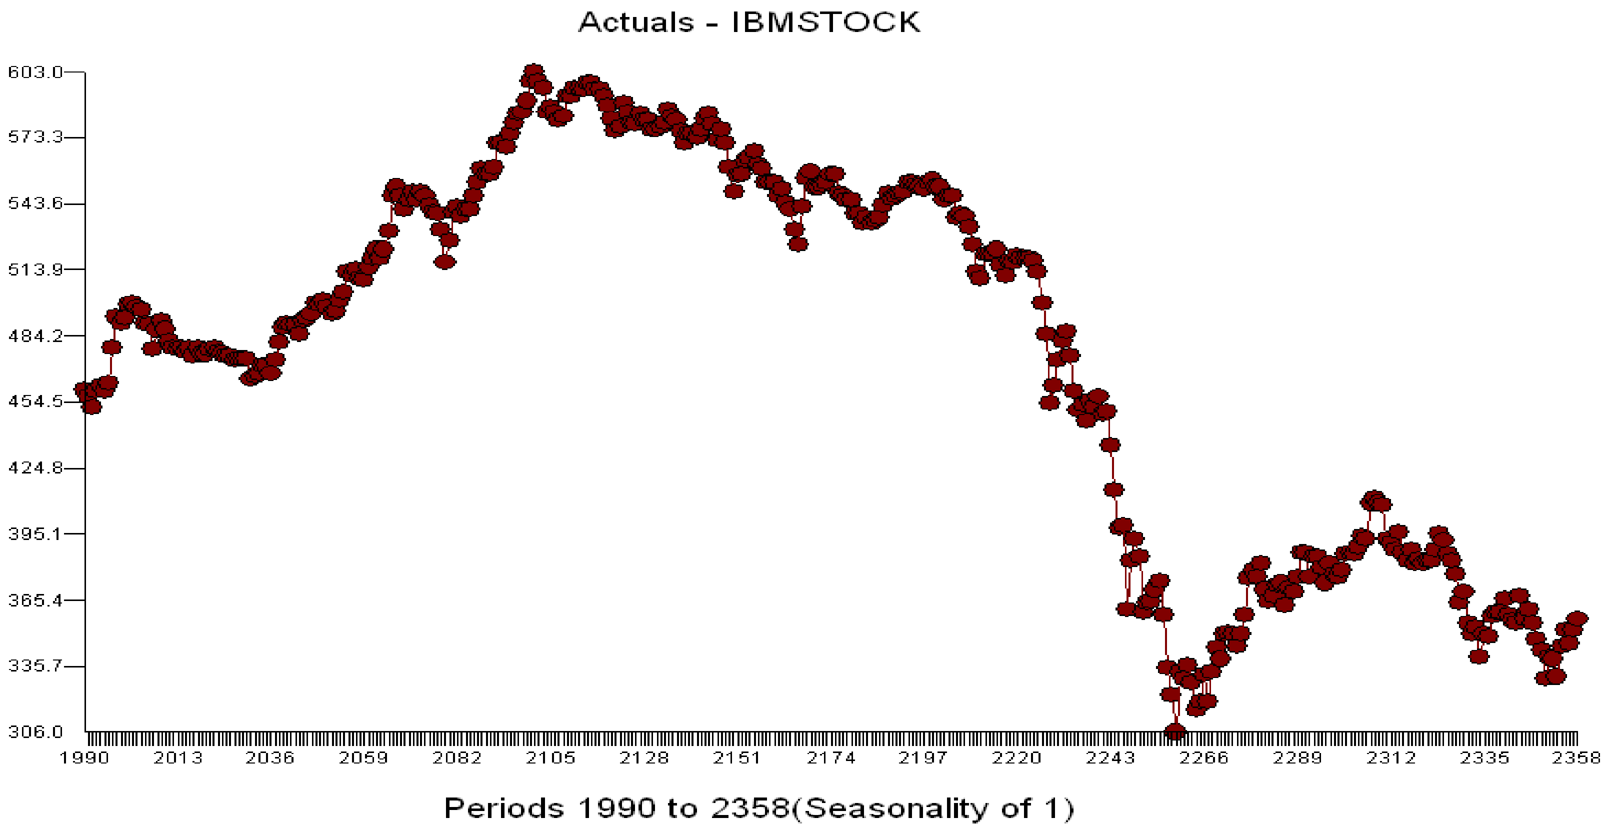

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Different Answers from Different Software for Same Problem

- While you might think that there is a standard answer for a standard model using the same dataset there isn't.   It might be due to mathematical inaccuracies in the computation.

- It is discussed by Yurkewicz in the OR/MS Today survey (see http://www.lionhrtpub.com/orms/orms-6-08/frsurvey.html for more - Forecast Pro, NCSS, Statgraphics. Systat, Minitab)

- Econometric Software Reliability and Nonlinear Estimation In Eviews, *Journal of Applied Econometrics*, Vol. 15, No. 1 (Jan. - Feb., 2000), pp. 107-110

- B. D. McCullough and Berry Wilson
"On the Accuracy of Statistical Procedures in Microsoft Excel 2000 and Excel XP," *Computational Statistics and Data Analysis* 40(4), 713-721, 2002

- B. D. McCullough and Berry Wilson
"On the Accuracy of Statistical Procedures in Microsoft Excel 2003," *Computational Statistics and Data Analysis* 49(4), 1244-1252, 2005

- B. D. McCullough and David A. Heiser
"On the Accuracy of Statistical Procedures in Microsoft Excel 2007" *Computational Statistics and Data Analysis* 52(10), 4570-4578, 2008

- The American Statistician
A Review of JMP 4.03 With Special Attention to its Numerical Accuracy by Micah Altman Vol. 15 No. 1 Feb, 2002, pp. 72-75

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Caveat Emptor

- Is the methodology explained? Does it do a “Pick Best” and ignore the assumptions?
- Does it build its own model from the information in each dataset or does it fit a couple of types of models to the data?
- Do they produce residuals that are free of pattern?
- Does the procedure explain what types of interventions it can detect?
- Does it explain what it does with interventions?
    - Does it correct for the outliers or just report them to you?
- Big ERP systems like Oracle, SAP, Manugistics, i2 are not exempt from this scrutiny

MERCURY MARINE

AUTOMATIC FORECASTING SYSTEMS

Autobox Support and Contact Information

Automatic Forecasting Systems, Inc. (AFS)

P.O. Box 563

Hatboro, PA 19040

Phone: 215-675-0652

Fax: 215-672-2534

email: sales@Autobox.com

Web Site: www.Autobox.com