ANALYSIS OF FLUIDITY INDICATORS AND SEASONALITY ADJUSTMENTS FOR CONTAINERS TRANSIT TIMES IN A MULTI-MODAL SUPPLY CHAIN

P. Boily and Y. Huang

Outline

I. Problem Description

II. Methodology

III. Illustration

IV. Consulting Post-Mortem

V. Additional Notes

PROBLEM DESCRIPTION

Problem Description

Supply chains play a crucial role in the transportation of goods from one part of the world to another. As the saying goes, “a given chain is only as strong as its weakest link” – in a multi-modal context, comparing the various transportation segments is far from an obvious endeavour.

Transport Canada is looking to produce an **index** to track container transit times in multi-modal chain networks.

This index should depict the **reliability** and the **variability** of transit times but in such a way as to be able to allow for performance comparison between differing time periods.

The seasonal variability of performance is relevant to supply chain monitoring and the ability to quantify and account for the severity of its impact on the data is thus of great interest.

Problem Description

The ultimate goal of this project is to compare quarterly and/or monthly performance data, irrespective of the transit season, in order to determine how well the network is performing, as it applies to the *Shanghai* ⟶ *Port Metro Vancouver/Prince Rupert* ⟶ *Toronto* corridors, and to produce a scoring methodology which could then be applied to other corridors.

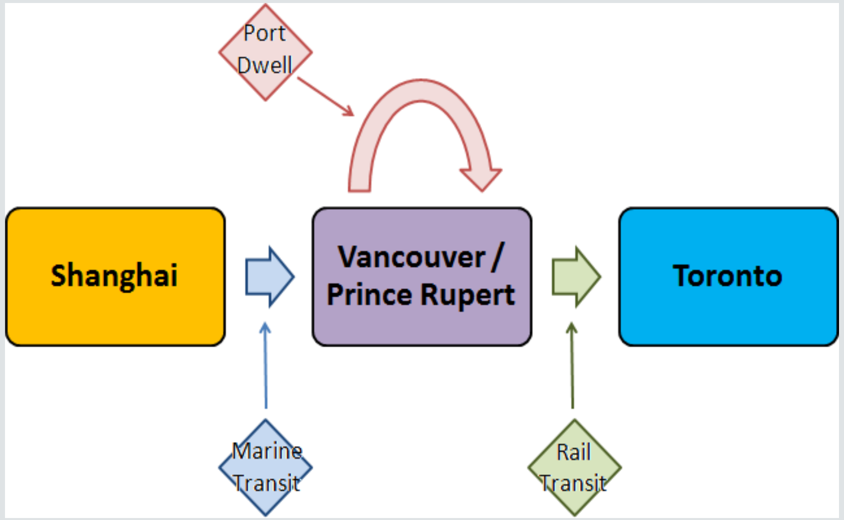

Problem Description

The supply chain under investigation has **Shanghai** as the point of origin of shipments, with **Toronto** as the final destination; the containers enter the country either through **Vancouver** or **Prince Rupert.**

Containers leave their point of origin by **boat**, arrive and **dwell** in either of the two ports before reaching their final destination by **rail.**

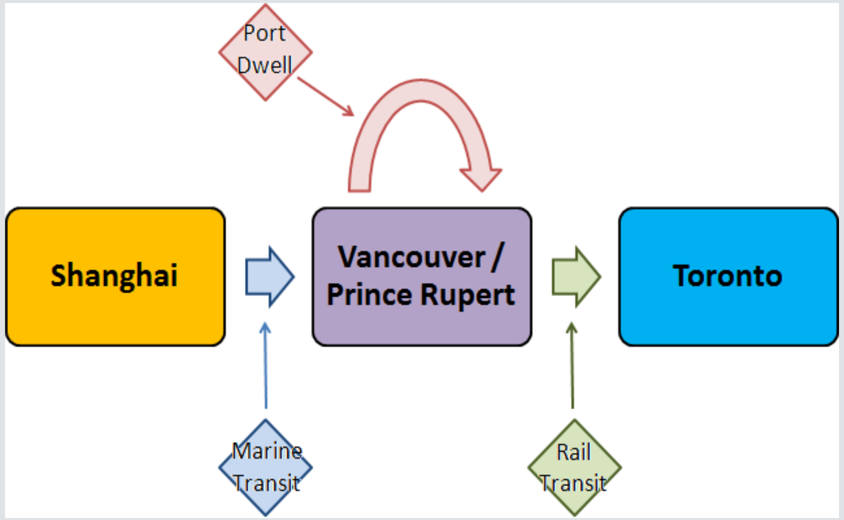

Problem Description

For each of the three segments (Marine Transit, Port Dwell, Rail Transit), the data consists of the monthly empirical distribution of transit times

- *from January 2010 to March 2013 (for Port Dwell)*
- *from January 2010 to April 2013 (for Marine and Rail)*

The data is built from sub-samples (assumed to be randomly selected and fully representative) of all containers entering the appropriate segment.

Problem Description

Each segment's performance was measured using Fluidity Indicators, which are computed using various statistics of the transit/dwelling time distributions for each of the supply chain segments. The main indicators under consideration were:

- ***Reliability Indicator*** *(RI) – the ratio of the 95th percentile to the 5th percentile of transit/dwelling times (a high RI indicates high volatility, whereas a low RI (≈ 1) indicates a reliable corridor);*
- ***Buffer Index*** *(BI) – the ratio of the positive difference between the 95th percentile and the mean, to the mean. A small BI (≈ 0) indicates that the mean and the 95th percentile transit times are roughly the same, and so that there is only slight variability in the upper (longer) transit/dwelling times; a large BI indicates that the variability of the longer transit/dwelling times is high, and that outliers might be found in that domain;*
- ***Coefficient of Variation*** *(CV) – the ratio of the standard deviation of transit/dwelling times to the mean transit time.*

Problem Description

The time series of monthly indicators (which are derived from the monthly transit/dwelling time distributions in each segment) were then **decomposed** into their

- *trend;*
- *seasonal component (seasonality, trading-day, moving-holiday), and*
- *irregular component.*

The trend and the seasonal components provide the **expected behaviour** of the indicator time series; the irregular component arose as a consequence of supply chain **volatility.**

A high irregular component at a given time point indicates a poor performance against expectations for that month.

ILLUSTRATION

Time-Series Decomposition

Basic Concepts

A **time series** is a sequence of values, measured at regular intervals over time. Ideally,

- *the reporting periods should be identical (e.g. daily, monthly, quarterly or yearly);*
- *the measurements should be taken over discrete (exclusive), consecutive periods,*
- *the concepts and the measurement approach should be consistent over time.*

Correlations and root causes can be identified if the time series are **stationary** (independent of time). Statistical tools which assume data independence are invalidated if the data is **serially dependent.**

If serial dependence is suspected or expected to exist, **time series decomposition** is required to identify **trend**, **cyclical** and **seasonal** components, in addition to providing accurate forecasts.

The time series seasonal adjustment enables the identification of **turning points** and provides consistent comparisons of indicators across time periods.

Components Hierarchy

Time series data can be broken down as follows:

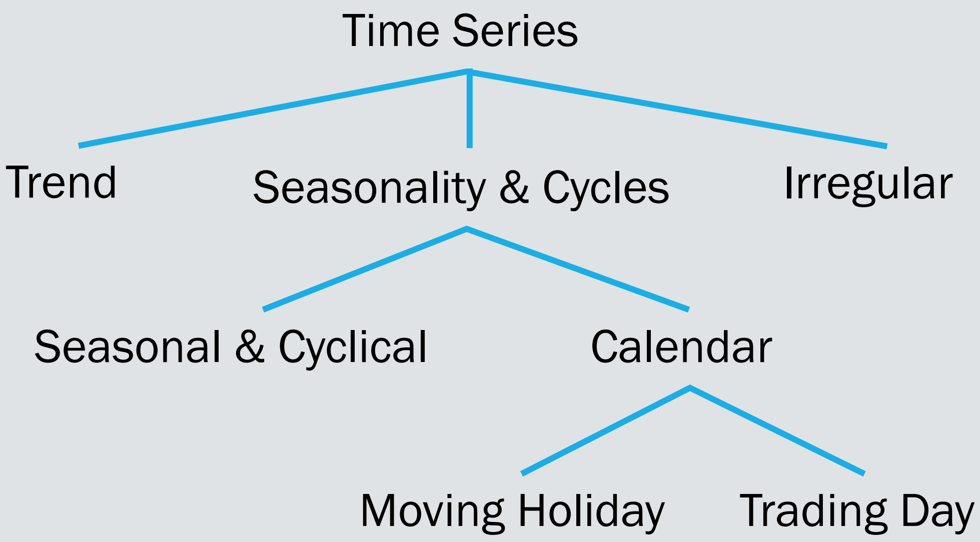

Breaks

Before carrying out seasonal adjustment, it is important to identify and pre-adjust for **structural breaks** (using the Chow test, for instance), as their presence can give rise to severe distortions in the estimation of the Trend and Seasonal effects.

- ***Seasonal breaks*** *occur when the usual seasonal activity level of a particular time reporting unit changes in subsequent years.*
- ***Trend breaks*** *occurs when the trend in a data series is lowered or raised for a prolonged period, either temporarily or permanently.*

Sources of these breaks may come from changes in government policies, strike actions, exceptional events, inclement weather, etc.

Models

Traditionally, the decomposition follows one of three models:

- *Multiplicative*
- *Additive*
- *Pseudo-Additive*

The choice of a model is driven by data behaviour and choice of assumptions.

The X12 model automates some of the aspects of the decomposition, but manual intervention and diagnostics are still required. X12 is implemented in SAS and R, among other platforms. **Consult the references for more information.**

Models – Multiplicative

This modeling approach assumes that

- *the magnitude of the seasonal spikes/troughs increases when the trend increases (and vice-versa);*
- *the **trend** T_t has the same dimensions as the **original series** O_t, and the **seasonal component** S_t and the **irregular component** I_t are dimensionless and centered around 1;*
- *the **seasonal fluctuation** $\sum_{j=1}^{n} S_{t+j} = n$, where $n = 365$ for daily series, $n = 12$ for monthly series, $n = 4$ for quarterly series, etc., and*
- *the original series O_t **does not contain zero values.***

Mathematically, the model is expressed as

$$O_t = T_t \times S_t \times D_t \times I_t,$$

where D_t is the **trading day component** due to **calendar effects.** All components share units.

Models – Multiplicative

After seasonality adjustment, the **seasonality adjusted series** is

$$SA_t = \frac{O_t}{S_t \times D_t} = T_t \times I_t.$$

After a log transformation, the **multiplicative model** becomes an **additive model:**

$$\log O_t = \log T_t + \log S_t + \log D_t + \log I_t$$

Models – Additive

This modeling approach assumes that

- *the* ***seasonal component*** S_t *and the* ***irregular component*** I_t *are independent of the* ***trend behaviour*** T_t*;*
- *the* ***seasonal component*** S_t *remains stable from year to year, and*
- *the* ***seasonal fluctuation*** $\sum_{j=1}^{n} S_{t+j} = 0$, *where* $n = 365$ *for daily series,* $n = 12$ *for monthly series,* $n = 4$ *for quarterly series, etc.*

Mathematically, the model is expressed as

$$O_t = T_t + S_t + D_t + I_t.$$

where D_t is the **trading day component** due to **calendar effects.** All components share units and dimensions.

After seasonality adjustment, the **seasonality adjusted series** is

$$SA_t = O_t - S_t - D_t = T_t + I_t.$$

Models – Pseudo-Additive

This approach assumes that some of the values of the original series O_t are 0 and that

- *the* ***seasonal component*** S_t *and the* ***irregular component*** I_t *are both dependent on the* ***trend level*** T_t*, but independent of each other, and*
- *the* ***trend*** T_t *has the same dimensions as the* ***original series*** O_t*, and the* ***seasonal component*** S_t *and the* ***irregular component*** I_t *are dimensionless and centered around 1.*

Mathematically, the model is expressed as

$$O_t = T_t + T_t \times (S_t - 1) + T_t \times (D_t - 1) + T_t \times (I_t - 1) \ = T_t \times (S_t + D_t + I_t - 2).$$

where D_t is the **trading day component** due to **calendar effects.** All components share units.

After seasonality adjustment, the **seasonality adjusted series** is

$$SA_t = O_t - T_t \times (S_t - 1) - T_t \times (D_t - 1) = T_t \times I_t.$$

Calendar Effects

A number of monthly and quarterly time series include calendar effects due to the varying lengths of the months, day-of-the week effects and holidays, fixed or moving:

- *The **trading day effect** is related to the monthly differences in the numbers of each day of the week from one year to the next (there may be more weekend sales in a month with five weekends). In each month, some days of the week may occur 5 times. The type of theses extra days affect the data for the month. Without an appropriate correction, it becomes impossible to compare monthly estimates from year to year.*
- ***Easter** is a **moving holiday** which occurs either in March or April between March 22 and April 25, or in the first or second quarter. Industrial production is lower in months in which Easter falls due to fewer working days. As a result, it needs to be removed before seasonality adjustment. If the Easter moving holiday effect is not corrected, the peaks and troughs due to its effects will be reflected in the final seasonal adjusted series.*
- ***Canada Day** is a **fixed holiday** that falls on each year on July 1 (unless July 1 is a Sunday in which case Canada Day is held on July 2). As a consequence, there is always an additional day in the first week of July where businesses are closed and the overwhelming majority of Canadians have the day off, usually modifying their regular activities. An adjustment must be made when looking at daily rates instead of absolute monthly numbers, say.*

Seasonal/Cyclical Adjustment Methodology

1. Choice of **seasonal decomposition model**
   - *Includes checking the need for data transformation as well as the selection of a specific decomposition model (multiplicative, additive or pseudo-additive) based on the data.*
   - ***Graphical inspection:*** *the multiplicative model should be used when the time series plot in continuous years shows that the size of the seasonal peaks and troughs changes as the trend changes; otherwise, the additive model should be used.*
   - ***AICC Comparison:*** *the log transformation should be selected if the Akaike Information Criterion* $(AICC)$ *satisfies* $AICC_{\mathrm{nolog}} - AICC_{\mathrm{log}} > 2$ *(these values can be computed using SAS's proc X12, for instance).*

Seasonal/Cyclical Adjustment Methodology

2. Adjusting and testing for **trading-day effects** (if applicable)
   - *Series should be adjusted to remove trading day effects from the final seasonally adjusted series when they are found to be statistically significant (unless the results are counterintuitive and do not match realistic day-to-day operations).*
   - *Testing for these effects includes looking for peaks in* ***spectral plots*** *(either the first differences of the adjusted time series or in the final irregular component, both adjusted for extreme values), the* $\boldsymbol{t}$ **−test** *or the* χ^2 **−test** *(significance of trading day effect regressors)*
   - *Significant* ***trading-day regressors*** *should be included in the model (X12 documentation).*

Seasonal/Cyclical Adjustment Methodology

3. Adjusting and testing for **trading day effects** (continued)
   - *There are instances when a series should NOT be adjusted for trading day effects: if the data-recording year is not divided in the same way as the calendar period described by months; if the data is collected at a point in time rather than (every day, say), or if the data is not collected in strict calendar months.*

4. Adjusting and testing for **moving holiday effects** (if applicable)
   - *Moving-holiday effects are identified using an* ***AICC test*** *(picking the model which gives the smallest AICC value for various moving holiday regressors) or* ***graphical inspection*** *(without adjustment for moving-holidays effects, the Seasonal Irregular component ratios or the month-to-month/quarter-to-quarter percentage change in the original series will not be consistent from year to year).*
   - *Significant* ***moving-holiday regressors*** *should be included in the model (more details are available in the X12 documentation).*

Seasonal/Cyclical Adjustment Methodology

5. Identifying and adjusting for **trend level shifts**
   - *Level shifts are abrupt but sustained changes in the underlying level of the time series associated with an unchanged seasonal pattern.*
   - *Without accounting for trend level shifts, there may be an **increased level of irregularity** for the seasonally adjusted series around the level shifts, increasing the volatility of the seasonally adjusted series.*
   - *These shifts can be identified using month-to-month/quarter-to-quarter percentage changes in original and seasonally adjusted estimates according to the following criterion: a sudden large increase which is not followed by a corresponding decrease, or vice-versa.*

Seasonal/Cyclical Adjustment Methodology

6. Identifying and adjusting for **outliers**
   - *Outliers are extreme values that fall outside of the general pattern of the trend and seasonal components, which can be caused by an extreme random effect or an identifiable reason.*
   - *Outliers found at the end of the time series may have a large impact on revisions when new data become available.*
   - *These can be identified using the table of final weights for the irregular component and the table of residual patterns of the irregular component, according to the following criterion: one or more of the values are 0.*

Data Quality Issues

Known data quality issues could affect the results of time series analyses:

- the method of data collection may lead to unusual effects, especially if collection is made on a non-calendar basis or if there is a lag between activity and measurement;
- any change to the method or timing of data collection could lead to the false identification of trend or seasonal breaks;
- some series are sensitive to events such as extreme weather, strikes, wars, etc., which could cause breaks or outliers of large magnitude;
- at least 5 years worth of data are required to insure stability on future updates, and
- at least 10 years worth of data are required to insure that the adjustment of the first year is unlikely to be revised.

ILLUSTRATION

Shanghai → Vancouver Marine Transit

Illustration

Shanghai → Vancouver

Let us illustrate the process of decomposition with the time series recording the **CV fluidity index** for the Marine Transit between Shanghai and Vancouver; the values are shown in the year-by-year plot to the right.

What trend(s) can be found in this data (if any)?

Are there months where the value of CV is “unexpected”?

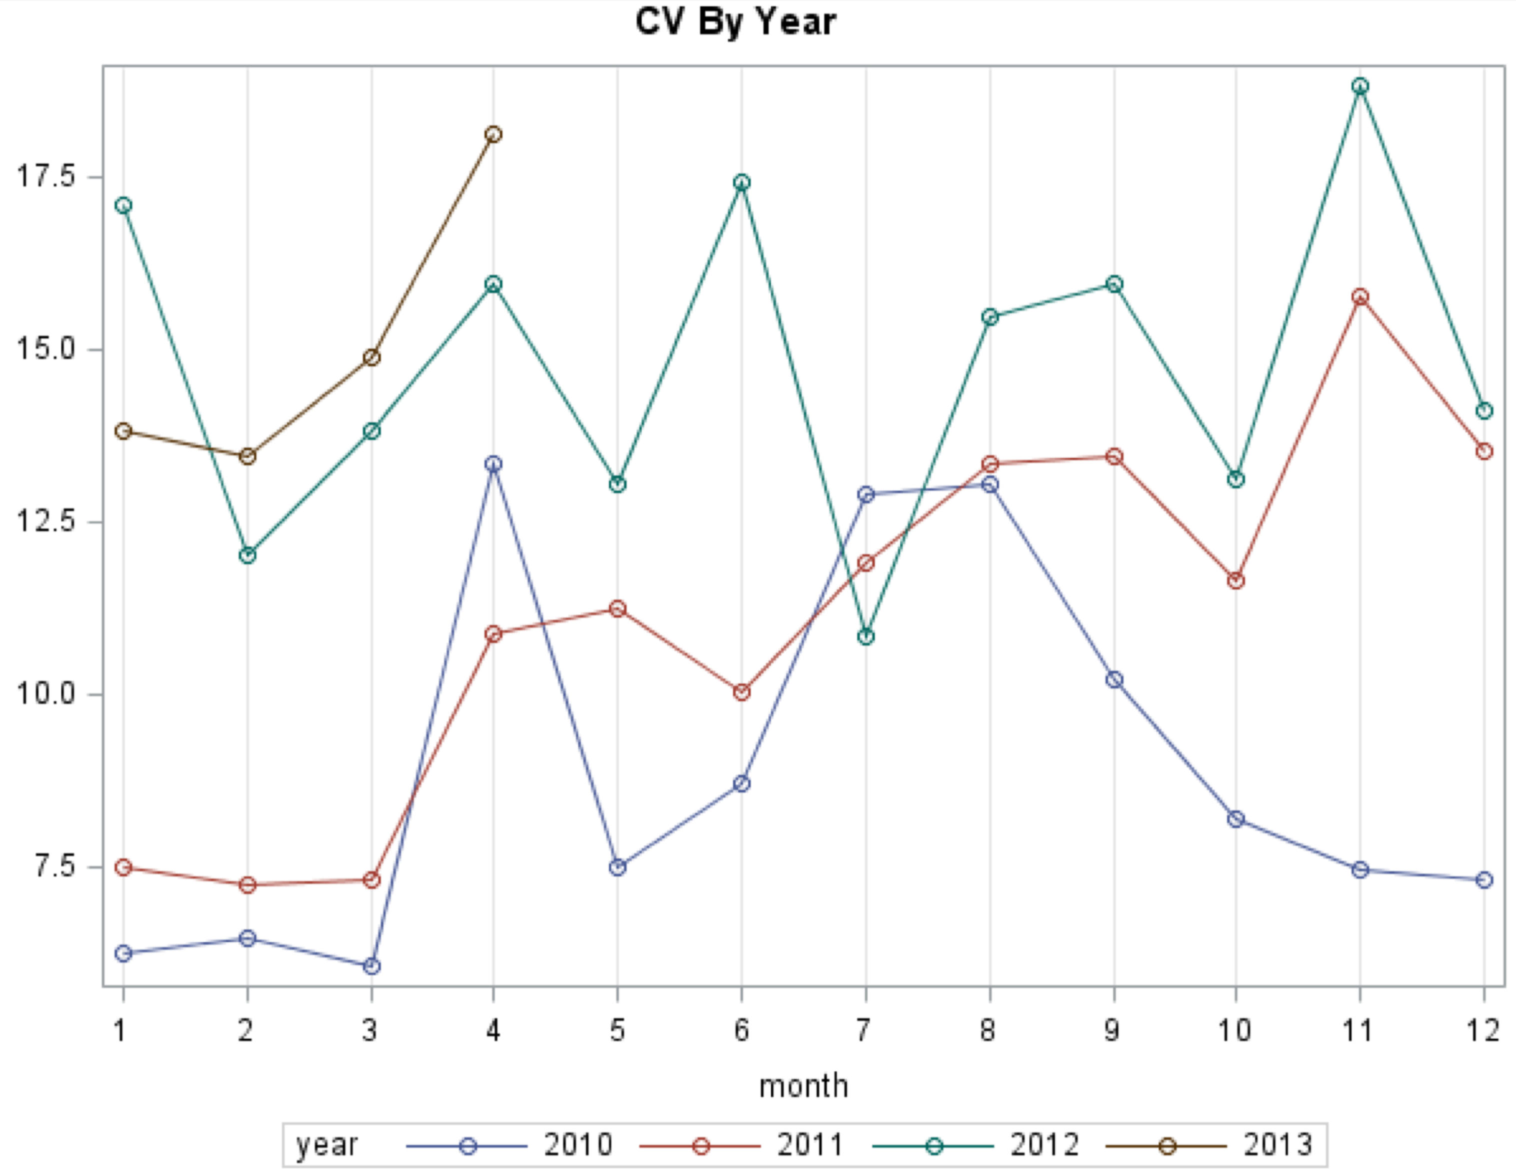

Model Selection

The continuous plot shows that the size of the peaks and troughs does not seem to change with changing trends: the **additive model** is thus selected.

SAS' `proc X12` agrees, and suggests no further data transformation.

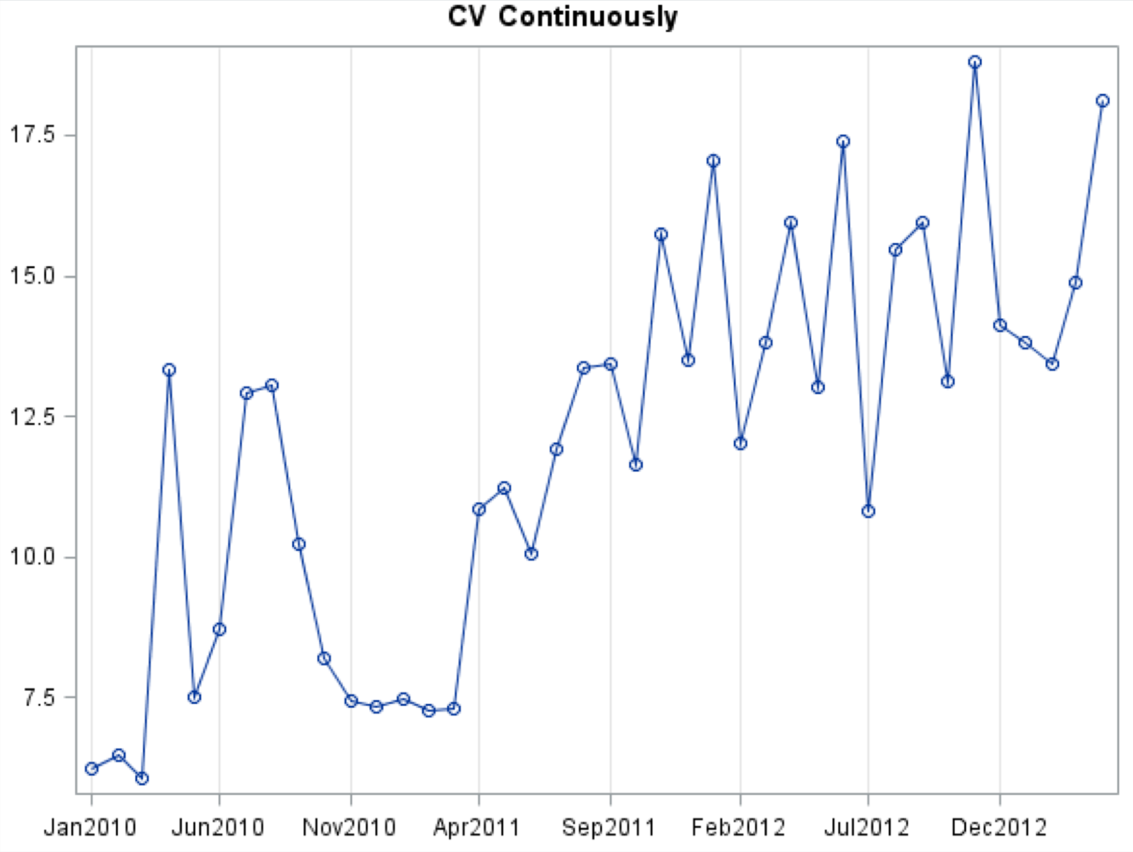

Estimation Summary For Variable hours_CV	
Number of Observations	40
Number of Residuals	27
Number of Parameters Estimated	3
Variance Estimate	5.6E-02
Standard Error Estimate	2.4E-01
Standard Error of Variance	1.5E-02
Log likelihood	0.4658
Transformation Adjustment	-69.5685
Adjusted Log likelihood	-69.1027
AIC	144.2053
AICC (F-corrected-AIC)	145.2488
Hannan Quinn	145.3613
BIC	148.0928

Results of Automatic Transformation Selection For Variable hours_CV	
AICC (with aicdiff=-2.00) prefers	No transformation
Adjustment will be	Additive

Trading-Day and Easter Effects, Level Shifts and Outliers

A spectral plot (not included) suggests that no **trading-day** effects are found.

At default critical values, the SAS procedure X12 further identifies an **Easter** effect but no **leap year** effect, as well as a suspected **level shift** in October of 2010.

Outliers were not detected.

Regression Model Parameter Estimates						
For Variable hours_CV						
Type	Parameter	NoEst	Estimate	Standard Error	t Value	Pr > |t|
Trading Day	MON	Est	-0.16078	0.30288	-0.53	0.5995
	TUE	Est	0.11693	0.31788	0.37	0.7156
	WED	Est	-0.41164	0.35394	-1.16	0.2540
	THU	Est	-0.88004	0.34186	-2.57	0.0152
	FRI	Est	2.47412	0.32547	7.60	<.0001
	SAT	Est	-1.64222	0.39435	-4.16	0.0002
	SUN(derived)*	Est	0.50363	0.37410	1.35	0.1883
Leap Year	Leap Year	Est	0.24294	0.97110	0.25	0.8042
Easter	Easter[8]	Est	-2.04419	0.94987	-2.15	0.0396
Automatically Identified	LS OCT2010	Est	-6.21498	1.34405	-4.62	<.0001

Diagnostic Plots

The **diagnostic plots** are shown on the next few slides: the 2010 CV series is prior-adjusted from the beginning until OCT2010 after the detection of a **level shift.**

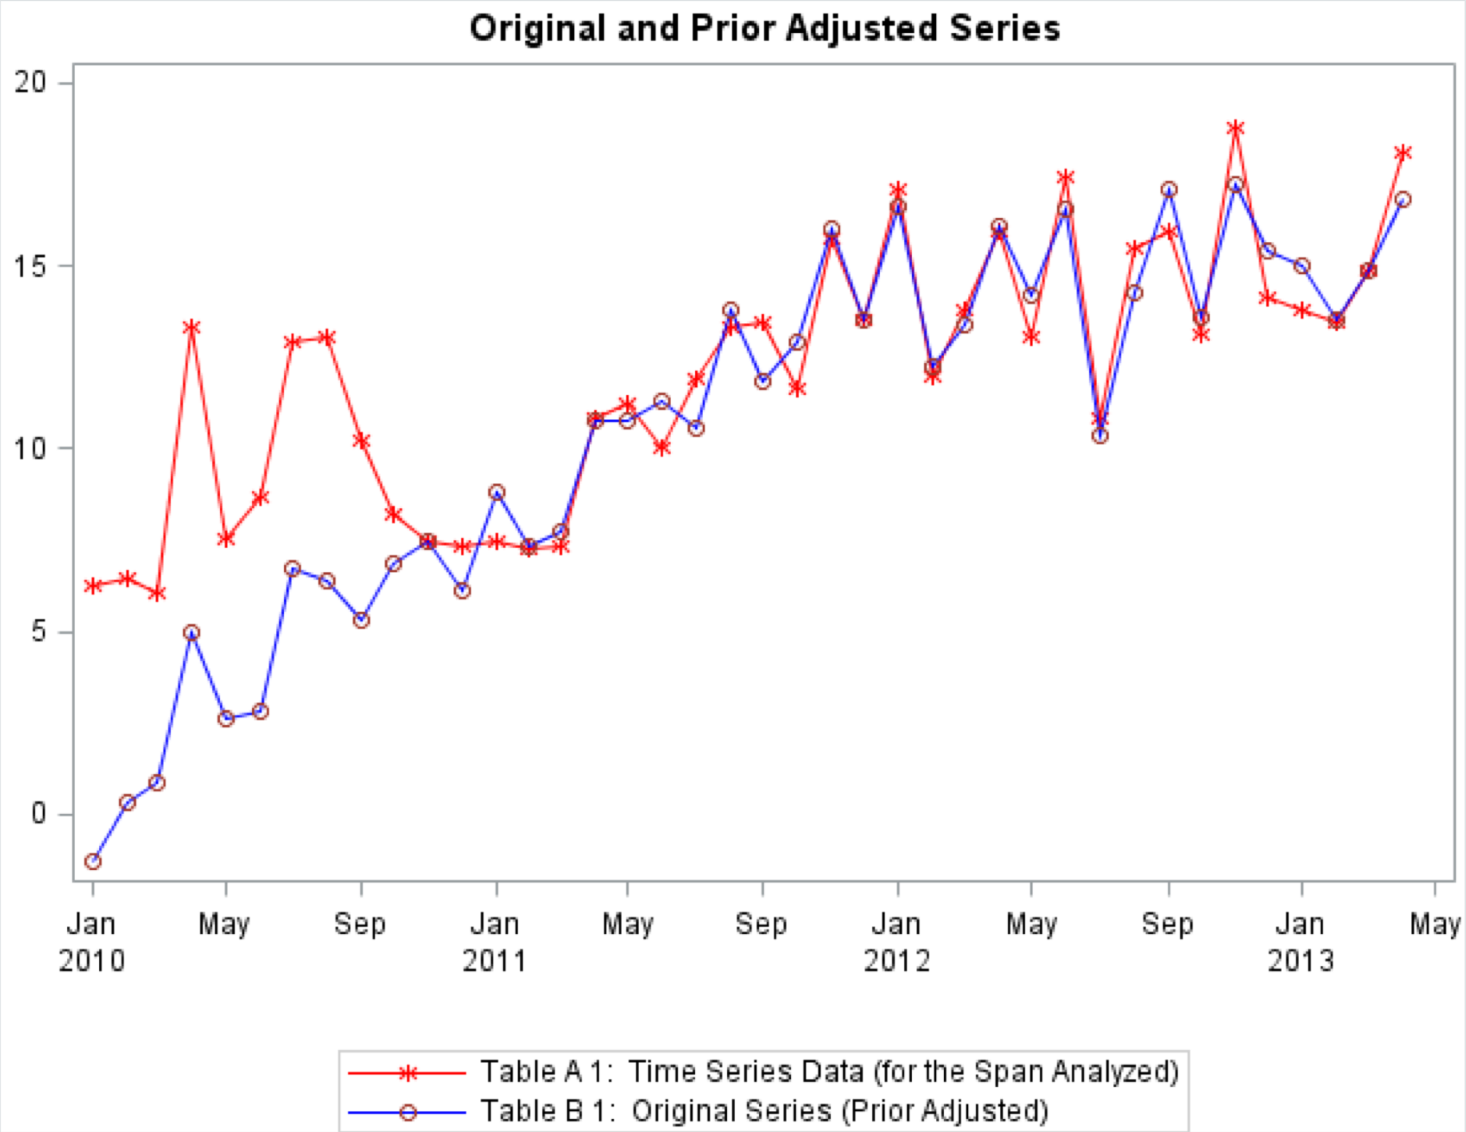

Diagnostic Plots

The SI chart shows that there are more than one **irregular component** which exhibits volatility.

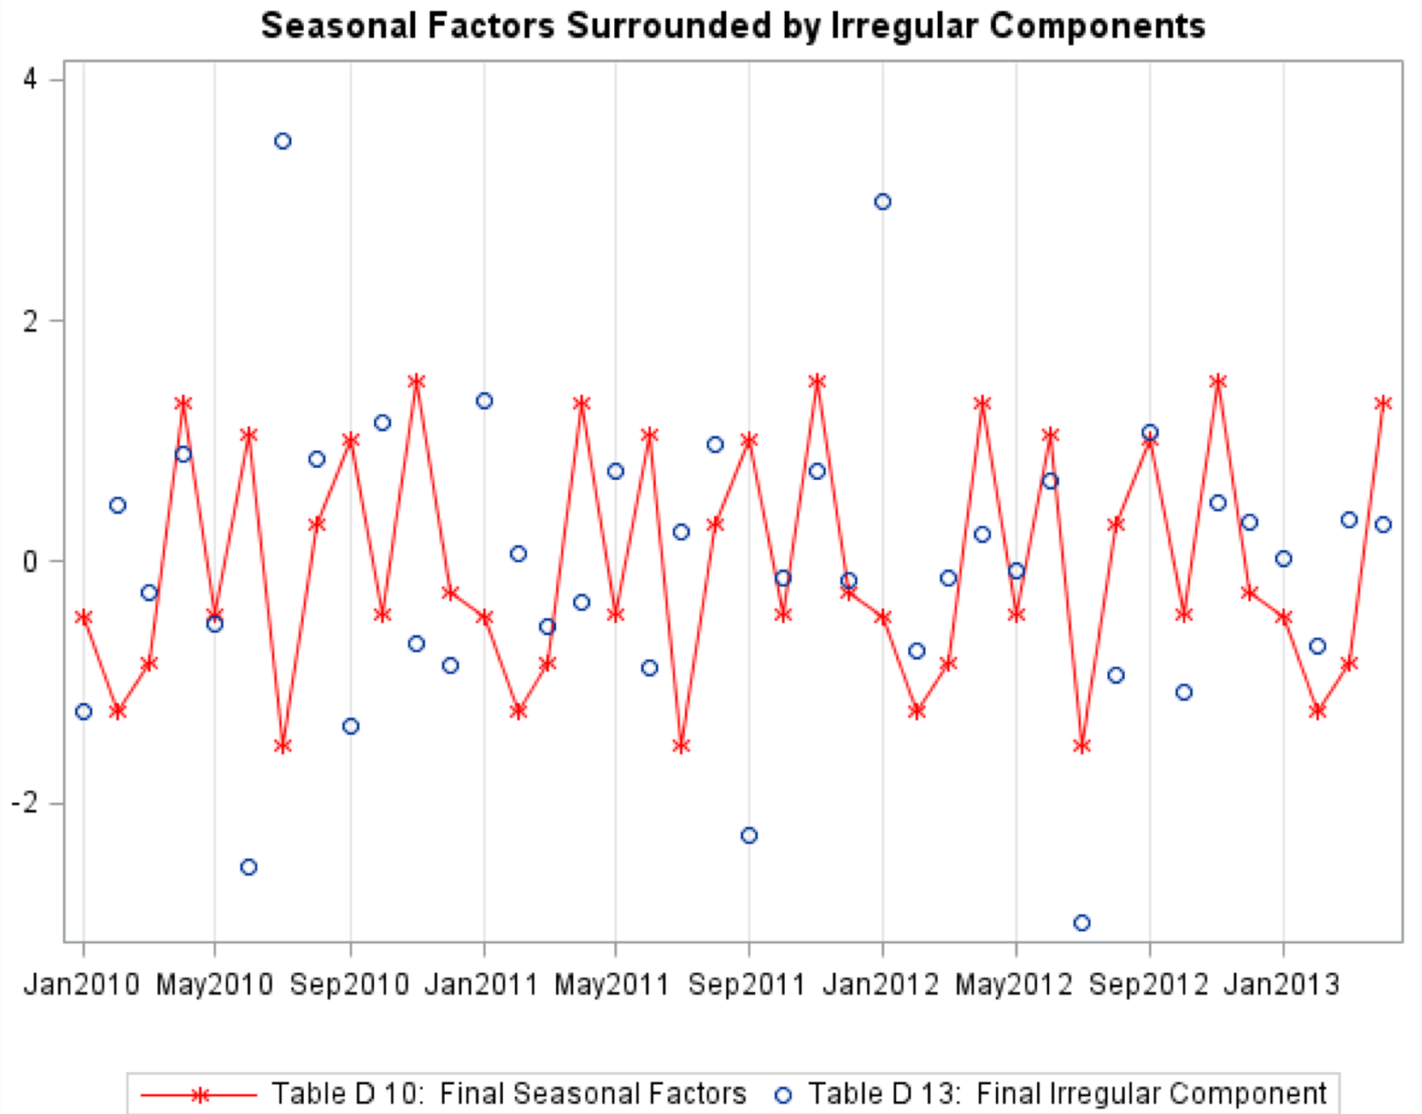

Diagnostic Plots

The **adjusted time series** is shown here.

Note the shift.

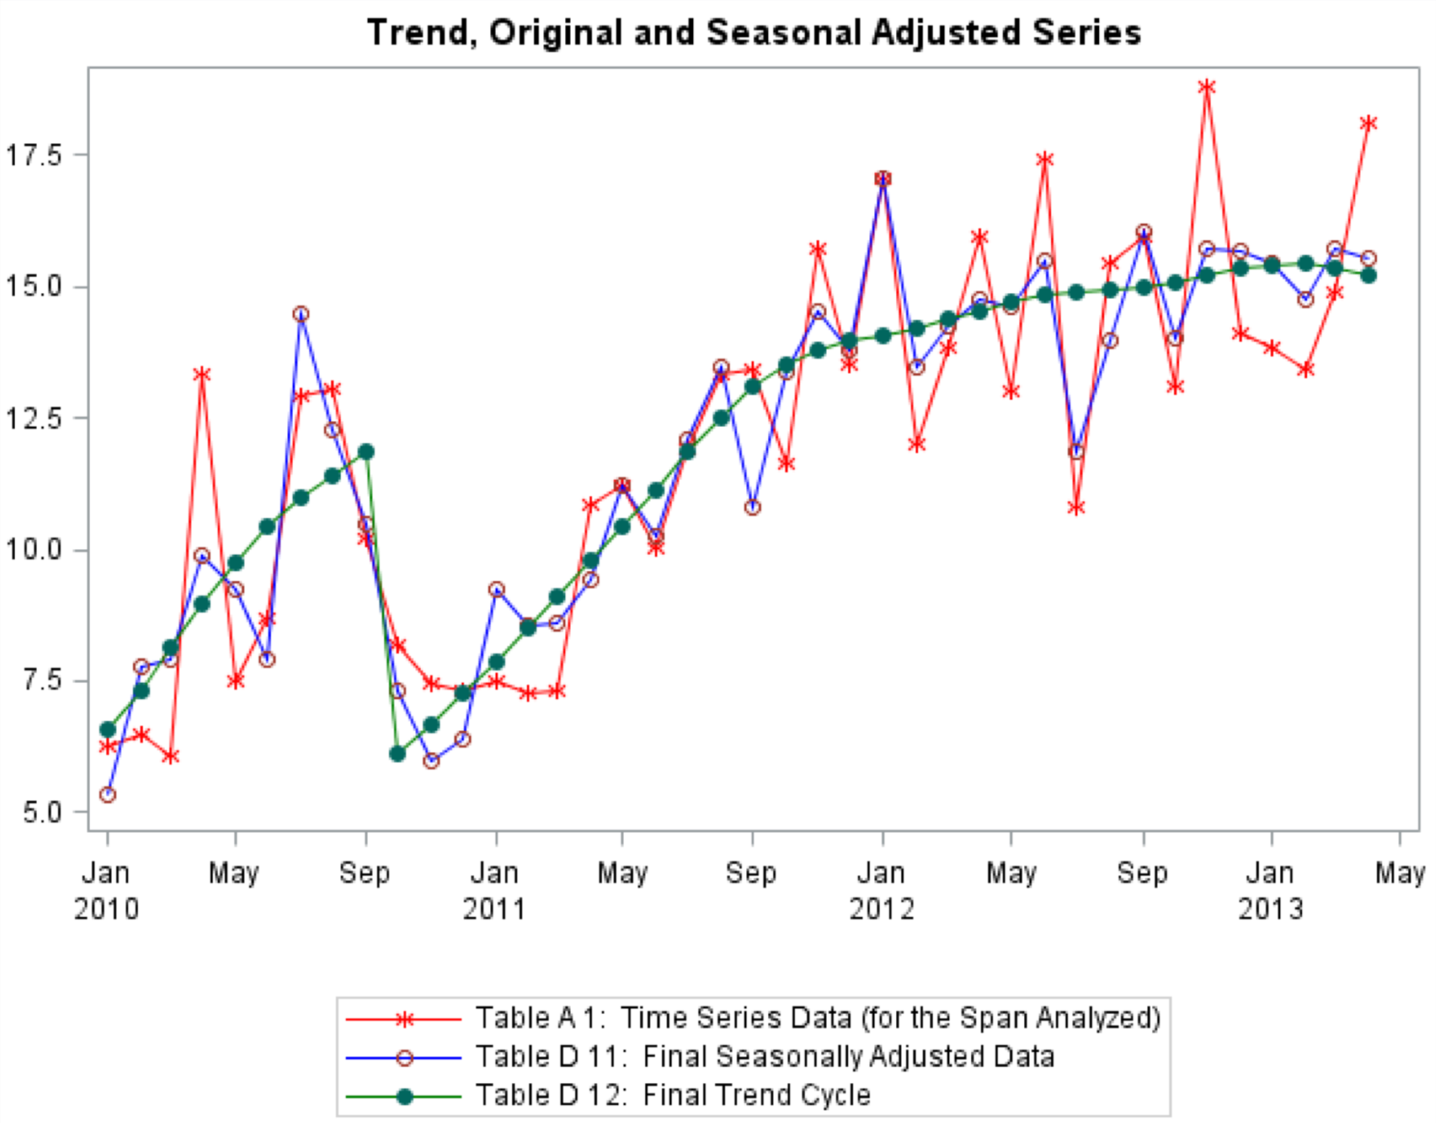

Diagnostic Plots

Now the trend and irregular components are shown separately.

What months should be audited?

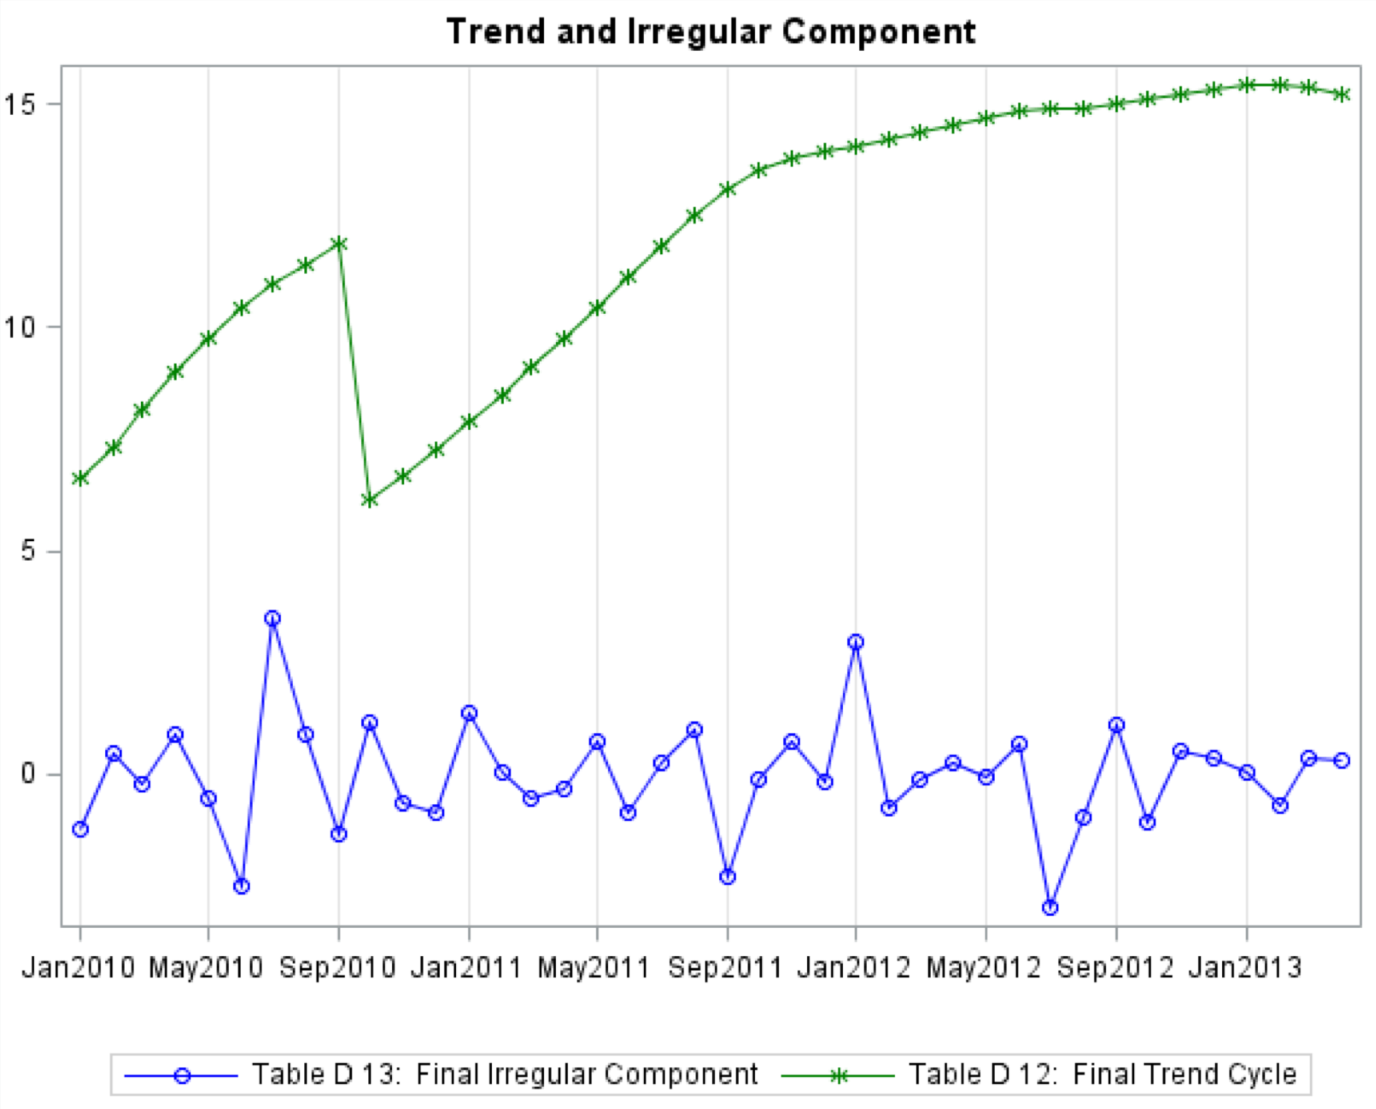

Comparison (Vancouver to Toronto BI)

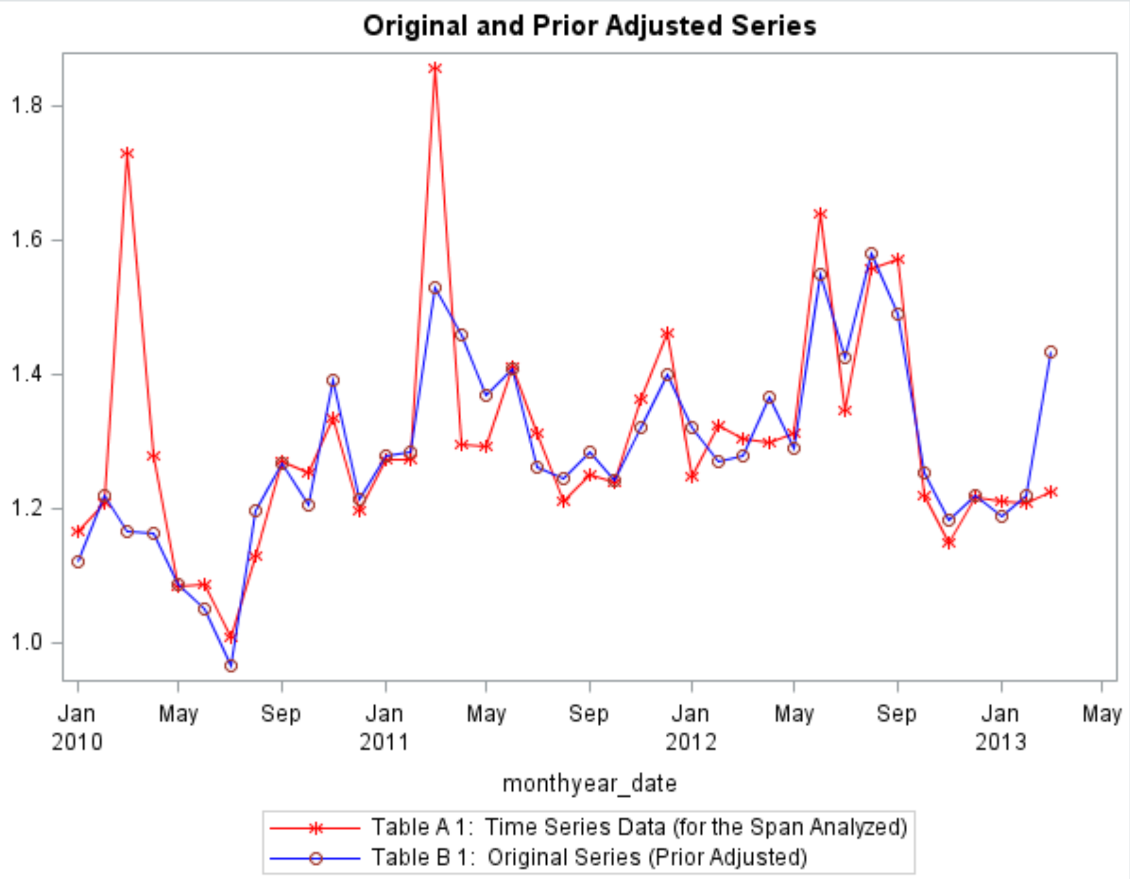

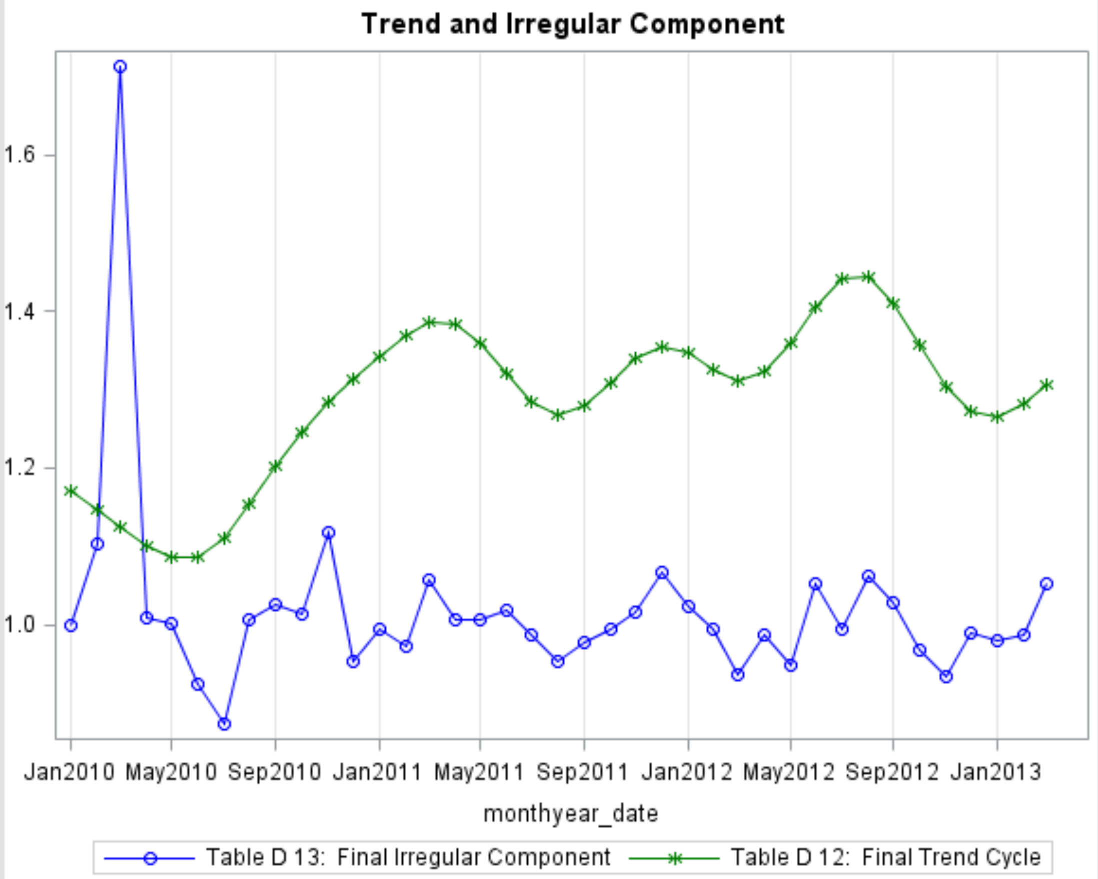

CONSULTING POST-MORTEM

Shanghai → Vancouver Marine Transit

Consulting Post-Mortem

Not enough data available (would have needed 5+ years), but OK since the focus was on the **methodology**; with data slated to come in indefinitely, the problem will eventually evaporate.

Transit dwelling time available only for a **sample of all containers** going through the supply chain; representativeness was questioned (no sampling design, inconsistent methods).

No overarching results that applied to all indicator time series, for each segment, except for the lack of a Chinese New Year effect, which was unexpected (given the origin of the chains).

Supply chain reliability is a function of the **total transit time** from its origin to its destination. End-to-end data was unavailable at the time.

Data security issues meant that consultants could not bring the data outside of TC's offices.

The client asked for an executive summary of the report but still hired another consultant to explain the report and the code; we **failed to recognize that the client was not understanding**, in part because of our frustration at severely undercharging the project ($4,424.73 for 250 hours). That's not a valid excuse.

ADDITIONAL NOTES

Forecasting Models, Missing Data, Automated Trend Extraction

Basic Notions

*Primary focus of forecasting is to try to **predict the future** using available data (time series or other).*

Forecasts tend to be **wrong**: aggregated forecasts are usually more accurate.

Emphasis should not be placed on a single estimate (the **mean**): forecasts should also include the **standard deviation** and an **accuracy range**.

Accuracy usually decreases when the **prediction horizon** lies further into the future.

Time series forecasts require the isolation of various patterns in the data: **trend**, **seasonality**, **cycles**, **level shifts**, **irregular components** and **outliers**.

Basic Notions – Methods

Data does not need to exhibit periodicity or time series characteristics, in which case a regular regression model could be appropriate.

In the presence of structure time series noise, a **Fourier** transform can help identify the number of distinct cycles (as well as their respective frequencies).

Most of the other time series method require the series to be **stationary** (i.e. the expected value of the series stays constant over time). Such a series can be represented by a model of the form

$$y_t = \mu + \varepsilon_t$$

where $\varepsilon_t \sim \mathcal{D}(0, \sigma^2)$ for some distribution with mean 0 and variance σ^2.

Most time series forecasting methods assume stationarity: if the series is not stationary, it must first be **decomposed** (detrended, deseasonalized) into its constituents components.

De-trending Time Series Data

Identification of trend in time series is subjective because what appears to be a trend over a short time period may prove to simply be a **small fluctuation** which could form part of a cycle over the long-term horizon of the series.

Regression models of various complexity levels can be fitted (against time and/or auxiliary variables) to identify possible trends. At long horizons, polynomial response functions explode: if such models must be used, **we recommend using linear or quadratic response functions**, as slope and concavity might be the best we can hope to detect in light of the previous remark.

In combination with appropriate data transformations (e.g. **logarithm**, **square root**, **inverse**, **Box-Cox**, etc.), the low order regression models can achieve good results.

Fourier transforms can help identify potential trend and cycles (as well as their respective frequencies), so can a variety of statistical tests (like the **Mann-Kendall test**, for instance).

De-trending Time Series Data – Methods

There are 4 main approaches to de-trending:

- ***finite differences:*** *iterated differences between subsequent time series observations, which can remove polynomial trends; useful if exact shape of trend cannot be estimated; too high an order may introduce variance inflation; ignores the potential effect of any variable over the trend, save for the passage of time;*
- ***curve fitting:*** *regression against time itself, or more complicated models involving auxiliary variables; prior knowledge of the situation can be used to provide an acceptable model which naïve analysis of the data might not be able to suggest; simple regression models may be unrealistic;*
- ***filtering and smoothing:*** *various weighted averages of the time series data can be used to compute a filtered series; advantages and disadvantages discussed in the next slides; the trend component output of the X12 procedure on the CV time series in the first section is an example; an explicit functional form for the trend is unlikely to be found;*
- ***cubic splines:*** *a separate cubic polynomial is fit continuously to every sequence of three points in the series; the first and second derivatives are continuous at each point; a "spline parameter," which depends on the relative importance given to "smoothness" and "goodness-of-fit" of the curve, is required to specify the spline flexibility.*

Notation and Forecast Evaluation

Let $x_1, x_2, \cdots, x_n, \cdots$ be the **past values** of the time series. In order to make a forecast at time t, we need to know $x_t, x_{t-1}, \cdots, x_1$.

The **forecast** $y_{t,t+\tau}$ is the prediction for $x_{t+\tau}$ made at time t; we use the shorthand notation $y_{t,t+1} \coloneqq y_t$ for the **next step prediction** of x_{t+1} made at time t.

The **forecast error** e_t at time t is the difference between the forecast at time t and the actual value of the time series at time t :

- *For a **multiple step** forecast:* $e_t = y_{t-\tau,t} - x_t$
- *For a **next step** forecast:* $e_t = y_t - x_t$

The **mean absolute deviation** $(\text{MAD} = n^{-1} \sum_i |e_i|)$ and the **mean square error** $(\text{MSE} = n^{-1} \sum_i {e_i}^2)$ can be used to compare the relative forecasting merits of various models.

Filtering Methods

Moving average of order N: arithmetic average of the most recent N observations:

$$y_{t+1} = \frac{1}{N}\sum_{i=0}^{N-1} x_{t-i}$$

- $\mathrm{MA}(N)$ *provides* ***stable forecasts****; bad data (e.g. irregular points, bad stretches)* ***is eventually removed*** *from the prediction process*
- *requires saving a lot of past data points;* ***lags*** *behind the actual trend; ignores complex relationships in data*

Weighted moving average of order N: attaches importance to certain observations in the form of weights (recent observations could have more influence than older observations, for example):

$$y_{t+1} = \sum_{i=0}^{N-1} w_{t-i}x_{t-i}$$

- $\mathrm{W}MA(N)$ *may* ***reduce the lag*** *shown by $MA(N)$ but there is no obvious way to introduce a weighing scheme $\sum_i w_{t-i} = 1$.*

Filtering Methods

Exponential smoothing with parameter $\boldsymbol{\alpha}$: weighted moving average with declining weights for past data: $y_{t+1} = \alpha x_t + (1-\alpha) y_t$.

- *By iterating the above relation, we see that* $\mathrm{ES}(\alpha)$ ***carries the entire past history of the series,*** *without the need to save past data points.*
- *Small values of* α *produce* ***stable forecasts with low variability,*** *but they increase the lag.*

Double exponential smoothing with parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ (Holt's Method): requires separate smoothing for the slope and the intercept if a linear trend is present:

$$B_t = \alpha x_t + (1-\alpha)(B_{t-1} + M_{t-1}), \quad M_t = \beta(B_t - B_{t-1}) + (1-\beta) M_{t-1}$$

- $\mathrm{HM}(\alpha, \beta)$ *makes* ***multi-step predictions*** *which can be quickly revised:* $y_{t,t+\tau} = B_t + \tau M_t$.

Triple exponential smoothing (Winter's Method) also incorporates a smoothing factor for **seasonal factors:** $y_{t,t+\tau} = (B_t + \tau M_t) c_\tau$.

- *Winter's Method requires two complete cycles to provide initial estimates, and a third cycle for fine-tuning.*

Other Considerations

Methods that are too sophisticated can be unreliable over the long-term.

Another family of methods to consider: Box-Jenkins, which require substantial data history, use the correlation structure of the data (none of the filtering methods do) and can provide much-improved forecasts in some situations.

Bayesian inference and Monte-Carlo Markov Chains could also be used if we have some prior information/belief regarding the structure of the time series and the auxiliary variables, but there is some controversy regarding non-frequentist approaches.

Automated Trend Extraction

Ideally, **trend extraction should not be automated**. There are ways to program the methods (such as proc X12 in SAS in R) to automatically select the model (additive, multiplicative, etc.), to search for outliers and level shifts, etc., but the tests that are used are NOT perfect and visual examination is typically needed to confirm the procedure's decisions.

That being said, it is possible to provide an algorithm that one would expect to de-trend and de-seasonalize a time series most of the time (the caveat being that unless one verifies the results on a given time series, one cannot be sure that the assumptions built into the algorithms applied to that time series).

MATLAB has a time series module; some of the documentation gives suggestions as to how to automate trend extraction. But MATLAB is not a native time-series environment: SAS or R are preferred alternatives.

A possible algorithm can be prepared once specific time series are exhibited.