Methodology: Seasonal Adjustment of the Transportation Services Index Time Series Data

The data underlying the Transportation Services Index (TSI) are seasonally adjusted by the Bureau of Transportation Statistics (BTS) using the methodology discussed here. This seasonal adjustment enables consistent comparisons of data between time periods.

Origin

To seasonally adjust transportation services time-series data, BTS adopted X12-ARIMA, Release 0.2, which was created by the U.S. Department of Commerce, U.S. Census Bureau. X12-ARIMA grew out of X11-ARIMA, which in turn originated from Census X-11 software developed at the Census Bureau in the 1950s and 1960s. The basic approach of X-11 is undoubtedly the most widely used statistical method for seasonal adjustment.[1]

Isolating Seasonality

The basic model of this approach is to decompose the time series into three components: trend (including cyclic phenomena), seasonality, and irregular. By a series of iterative steps, the seasonality component is eventually isolated and removed from the original data series. In applying this methodology to the transportation services time-series data, we found that each element of the TSI—rail (passenger and freight), pipeline (petroleum and natural gas), transit, waterborne, trucking, and aviation (passenger and freight)—displays strong seasonality patterns. However, in some series the seasonality was less pronounced than in others. For example, transit seasonality must be isolated from a background of considerable fluctuations or noise. When there is a great deal of fluctuation in the data, as in this instance, the seasonality component is much smaller relative to the other two components (trend and irregular).

Holiday and Trading-Day Effects

The seasonal effect in a time series is any effect that is reasonably stable in terms of annual timing, direction, and magnitude. Seasonal adjustment is the process of estimating and removing the seasonal effects from a time series. Because the seasonal effects can disguise important features of economic series such as direction, turning points, and consistency between other economic indicators, seasonal adjustment can also be thought of as focused noise reduction (Hood, 2009; Ashley, 2001).

In addition to seasonal effects, monthly time series that are totals of daily economic activities are frequently influenced by the weekday composition of the month (Findley et al, 1998). Some months are longer than other months, which can affect the monthly totals of output services. Recurring effects associated with weekday composition in monthly (or quarterly) economic time series are called trading day effects.

Certain kinds of transportation services and their associated time series, such as aviation (passenger and freight), are affected significantly by holidays. Effects from holidays, such as Christmas, that always occur on the same date of a month each year are seasonal components of a time series. Effects associated with holidays that are not always on the same date of a month, such as Labor Day, Thanksgiving, and Easter, are called moving holiday effects.

Both trading day and moving holiday effects are persistent and predictable calendar-related effects. These calendar effects also need to be removed from the seasonally adjusted series. Not surprisingly, trading-day and moving holiday effects influence many of the time series used as components in the TSI.

Outlier Adjustments

In all of the transportation services time series it was necessary to make adjustments to the defaults offered in the software. One of the series used in the TSI (waterborne) required an adjustment to the software default 13-term Henderson moving average for the trend-cycle component estimation. Additionally, each of the series had too many extreme observations and outliers, when using the default lower sigma limit of 1.5 to estimate the seasonal component. Because the basic tool of both the X-11 and X-12 ARIMA are moving averages, which are linear operators, they respond dramatically to the presence of extreme observations (Ladiray and Quenneville, 2001). When doing seasonal adjustment, too many outliers may interfere with estimation of the seasonal component and make seasonal adjustment unstable. We had as our objective to ensure that no more than one out of eight data observations were of the extreme variety. The criterion of one out of eight comes from best practices in the field. For the most part we met that objective.

Two Kinds of Outliers: Extreme Values and Worst Outliers

We limited the number of extreme values by increasing the lower sigma limit to 1.8. This is justified with the transportation services time series data because the 10 transportation time-series data sets are strongly impacted by the vicissitudes of weather patterns as well as a host of unknown variables, making transportation time-series data inherently variable. Our experience has shown that nearly all time series need a similar adjustment to reduce extreme values.

Worst Outliers

Using a variety of procedures, including a hard look at changes in the monthly data, we detected several extreme outliers. An example of an extreme outlier that strongly impacted many of the time-series data is the terrorist attacks of September 11, 2001. The impact of 9/11 was particularly powerful on aviation passenger and freight outputs.

We intend in the future to hard-code the most extreme observations as outliers when we have external information (e.g., catastrophe, strike, weather) that points to their origin. A few extreme outliers found by the program series are noted in the data summaries that follow.

Extreme values and outliers—whether found by the X-12 program or other sources—are adjusted out of the data series when estimating the seasonal component so that they do not affect the estimate. These outliers and extreme values are included with the irregular component. Because the seasonally adjusted series is the trend and irregular components together, all outliers and extreme values are included in the seasonally adjusted series.

Modeling the Time Series: Multiplicative or Additive?

A series is modeled in one of two forms: multiplicative or additive. In practice, when the peaks and troughs remain constant as the trend increases or decreases, the additive model is used. When the peaks and troughs expand as the trend increases or decreases, the multiplicative model will fit the time series better. In the seasonality adjustment for the June 2004 index, we determined the best model using the Census Bureau's X12-ARIMA monthly seasonal adjustment method. Previous adjustments had used the multiplicative model. However, it was found that five of the series were best rendered by an additive model: trucking, transit, rail freight (carloads and intermodals), and waterborne. In these instances, the model adopted was changed to "additive" for the duration of the TSI's experimental phase.

Moving Seasonality

Adjustments were made to the seasonally adjusted data based on several criteria developed to assess the basic model, which are given as measures M1 through M11. These criteria involve the way the trend line is estimated, the amount and type of variation of the irregular component, and the stability of the model, especially indications of moving seasonality. The latter occurs when the high and low months and undulations in the data are shifting, thereby introducing additional variability. Although moving seasonality should be a matter of concern, its existence does not preclude executing seasonal adjustments.

Moving seasonality is present in the February 2009 time series for transit, aviation passenger, and natural gas (pipeline).

Meeting Criterion

Failing a single criterion by no means implies that the model is not good, just as passing all the criteria does not ensure a good adjustment. These measures only serve as aids to guide the adjustment. A criterion is met if it is less than 1.0; failure means that the measure was equal to or greater than 1.0. Gradually, with experience, we have learned to put the most emphasis on one measure, M7, which determines if there is any identifiable seasonality, and to be very indulgent with all the others. A failure in the other measures may mean, for example, a weak trend line, or a lot of autocorrelation, or shifting seasonality; nevertheless, although these other commonly computed measures may fail for some of the TSI series, they do not act as show-stoppers. By using M7 as the dominant criterion, all the series used in the TSI display very strong seasonality components.

Months to Cyclical Dominance (MCD)

One of the measures, M5, utilizes the concept of months to cyclical dominance (MCD). This is an indicator of how strong the trend is relative to the irregular component. MCD is defined as the number of months it takes before variations in the trend are expected to become larger than variations in the irregular component. When the irregular is under control, the MCD is about 6. An MCD of 6 would mean it takes approximately 6 months before changes in the seasonality adjusted series can be attributed to changes in the underlying trend rather than the irregular. For the TSI, the MCDs of the 10 individual time series range in value from 3 to 12. This time, pipeline gas reached an MCD of 12. Until April 2007, waterborne’s MCD was 12. An adjustment in the waterborne series changed that. The MCD value is noted in summaries of each series that follows this introduction.

The MCD is impacted by the Henderson moving average filter used to ascertain the trend. Except for waterborne, we have elected to use the default—a 13-term Henderson moving average. This means that the data are averaged in groups of 13 adjacent data points. A shorter length of the Henderson average would bring the MCD down. However, this procedure will make the trend line less smooth.

Standardizing Each Series

The nearly dozen time series used to create the TSI had varying amounts of historical data. Some series began as early as 1973. Earlier data may be useful for historical purposes, but are no help in seasonally adjusting data in recent years. Why? The answer is that the X12-ARIMA program doesn't allow the earlier data to carry much weight for the years at the end of the series. For this reason, with the exception of transit, all time series used in the current (2008) TSI begin at January 1990. This is more than sufficient time to obtain a good seasonal adjustment

Waterborne

Comparing 12-month periods in the series, it is clear that January and February are the low points. March through August, October, and November are the peak months, with July (slightly) the highest point. No substantial amount of moving seasonality was detected. Extreme outliers include January and June 1995, October 1997, and January and December 2000. The years 1995 and 2000 were more erratic than the other years, but no level shifts were detected. December 2006 and January 2007 were also extreme outliers.

The waterborne series is composed chiefly of three series: coal, petroleum, and farm products. These series are largely independent of each other. The coal series is not seasonal, while the farm products series is very unstable. After January 1994, the coal portion was left alone (i.e., not seasonally adjusted), the remainder was seasonally adjusted, and then coal was added back in. An additive model was used to fit the series. The effects of trading day and Labor Day were significant.

The months to cyclical dominance (MCD) value for waterborne denuded of the coal component was unduly high, 12, even without the coal to confound things. This meant that for the waterborne data without the coal factor, it would take a whole year for changes in the trend to equal changes in the irregular, indicating that the trend in waterborne is either very weak or the irregular component is very strong, or both. To compensate for this weak trend or large irregular component, a shorter averaging procedure was employed to better represent the true trend of waterborne. The Henderson moving average default of 13 was changed for waterborne to 7. This enabled the MCD to yield a value of 3.

NOTE: For the February 2009 TSI, no data values for waterborne or coal were required to be forecasted and used in the final model.

Rail

Freight. Two time series constitute rail freight: carloads and intermodals. For both time series, we find a double peak in August and October with October as the larger of the two. February is reliably the low month. The moving holiday effects were not significant in either series with the exception of Easter and Thanksgiving in intermodals. Notable outliers for carloads were August 1993, and January in 1996 and 1998. We haven't seen extreme outliers in carloads for the last six years. For intermodals, April 1991, June 1992, October 2002, and June through September 2004 were extreme outliers as was February 2008. Trading-day effects were a significant factor in both carloads and in intermodals in the June 2007 update. Rail freight is very seasonal; the adjustment in both series reveals clear trend lines and stable seasonality.

No significant moving seasonality was detected in either series. The months to cyclical dominance (MCD) values are 5 and 3 for carloads and intermodals, respectively. These MCD values are a little above and a little below, respectively, the median value of 4 of the 10 transportation series used in creating the TSI. Both series receive a "pass" on all diagnostics, M1 through M11.

Each of the rail time series was first seasonally adjusted and then combined into a rail output index. Additive models were both used to fit carloads and intermodals time series.

NOTE: For the February 2009 TSI, no data values in either series (carloads and intermodal) were needed to be forecasted and used in the model.

Passenger. Like freight data, rail passenger data are also well behaved with a clear trend and consistent and stable seasonality. Rail passenger adjusts similarly to rail freight-intermodals. Trading-day effects were not significant, but Easter week, Labor Day week, and Thanksgiving were determined significant factors in rail passenger. No significant moving seasonality was detected. The high points of the year are July and August; the low months were in the winter with February most often the lowest. Extreme outliers found in the series were April 1991, June 1992, and February 1998. Like rail freight, rail passenger easily passed the diagnostic standards, M1 through M11. Its MCD value of 4 is the median for the 10 TSI series. Rail passenger data are best modeled with the multiplicative model.

NOTE: For the February 2009 TSI, two data points were forecasted and used in the model for rail passenger.

Pipeline

Petroleum. Petroleum and petroleum products monthly amounts show a distinct pattern of February as consistently the low month. May, July, August, October, December, and January are the peak months, with August slightly more often the annual high month. Significant effects for trading days were detected, but no significant moving holiday effects were presented. The number of outliers was reduced by adjusting the weighting as in the other series. In January 1985, a new upward trend occurred. When seasonality is removed, the monthly increase (i.e., from December 1984 to January 1985) was 37 percent from the previous month. This level is sustained over the following months, never to return to its previous level. This change in the trend was one reason why we start the series a little later at January 1990. Extreme outliers include March 1991, February 1992, October 1994, and February 2000. February occurs as a moderate and extreme outlier quite frequently, indicating that the model is compromised to some degree by the relatively extreme results in February.

The criterion for each of the diagnostic measures, M1 through M11, was easily met for the Petroleum series. No significant amount of moving seasonality was detected. The MCD value was 6, which is above the average of 4. Multiplicative models were used for both series of pipeline.

NOTE: For the February 2009 TSI, two data points were forecasted and used in the model.

Natural gas. The peaks and low months in natural gas complement the petroleum time series. January, February, March, and December are the high gas consumption months each year. Gas consumption is relatively low in May, June, September, and October. The final adjustment included a change in the weighting to manage the number of extreme values. Like petroleum, natural gas data series is impacted by trading days but moving holiday effects were not significant. Three extreme outliers were February 1994, December 2000, and October 2001. The criteria for diagnostic measures M3, M4, and M5 of pipeline gas were not met. M3 measures the amount of month to month change in the irregular component compared to the amount of month to month change in the trend-cycle. M4 indicates autocorrelation in the irregular component of the model. M5 measures the number of months it takes the change in the trend-cycle to surpass the change in the irregular. M3 and M5 can both fail if the trend is flat. If this is the case, the seasonal adjustment is not necessarily “bad”; however, it will tend to be dominated by the irregular and thus movements in the seasonally adjusted series will not be reflected in movements in the trend. We now believe this is not a problem in the deseasonalized data and did not attempt to eliminate it. While strong seasonality is not in doubt, it is confounded by the existence of significant moving seasonality, which was detected at the 5 percent level. This means the seasonality tended to shift, adding an element of uncertainty to the adjustment. The irregular component is rather large in natural gas; the number of months to cyclical dominance (MCD) is 12, which is the highest of the 10 times series in the TSI. Because pipeline gas had an MCD value of 12, pipeline gas data, in general, shows strong irregularity components and/or weak underlying trends. Not surprisingly, the flow of gas or oil in the pipelines is not influenced by the holidays. However, Trading Days impacts their flow significantly.

NOTE: For the February 2009 TSI, no data points were forecasted and used in the model.

Trucking

October appears to be the most frequent peak month in the calendar year. Other high points are August, March and June. The low points occur during the late fall and winter months from November through February; February and December are most the two lowest months. The seasonality aspect is very strong and consistent throughout the decades. No evidence of significant moving seasonality was detected. Only three extreme outliers were found—April 1994, December 1994, and April 2000. At this time, no explanation has been found for the surprisingly large output for December 1994, where the irregular component was 7 standard deviations above the average of the series. The outliers appear to be concentrated in the year 1994, indicating that this year was not normal. The average monthly standard deviation for the year 1994 exceeded three standard deviations.

Currently, trucking is affected by trading-day effects and Easter week. In the past, Thanksgiving has been a significant factor, but this holiday doesn't appear to have any effect now. The irregular component is relatively small in trucking; the number of months to cyclical dominance (MCD) is 3. The deseasonality conducted on the trucking data met all diagnostic measures, M1 through M11. Since the November 2004 TSI, the model for trucking was changed to an additive model. The multiplicative model was more valid in the past because of the lower fluctuations of the series apparent during the earlier years. As the series has progressed over time, the fluctuations are becoming approximately constant as the trend increases.

NOTE: For the February 2009 TSI, no data values were forecasted and used in the model. The "advanced" value from the American Trucking Associations is used for the TSI computations, and it is finalized in the following month.

Transit

An examination of the raw data going back to 1979 shows evidence of seasonality. The lowest month in ridership is February. The high month is October. The pattern of seasonality is statistically significant, but the frequent shifts in the troughs and peaks make the seasonal adjustment a real challenge. Although the series is thus highly erratic, there is a little more stability in recent data. For this reason, we decided to work with just the last 13 years.

The day of the week plays a strong role, particularly Thursday through Sunday, with Sunday having an especially large impact. Surprisingly, most moving holiday effects were not worth hard-coding because the day of the week absorbed most of the moving holiday effects. The effect of Easter holiday was highly significant in the transit series. Significant moving seasonality (1% level) was detected in the series. Each of the diagnostic measures passed the criterion. Two extreme outliers were found in June: 1994 and 1998. The other two extreme outliers were December 1993 and April 1995. An additive model was used to fit transit time series data. The number of months to cyclical dominance (MCD) is 4, which is the median for the 10 series.

NOTE: The most recent data available for transit is December 2008. To make up for the data gap, two data points were required to be forecasted for the February 2009 TSI. Normally, Transit lags well behind the other series in up-to-date data. For example, in the December 2004 update, 8 years of transit data were revised.

Aviation

Freight. The seasonality pattern in aviation freight shows high points most often in October and low points in February. Since about 1986 there has been a steady increase in tons carried. As this trend has risen, the seasonal fluctuation has increased markedly. These data were relatively easy to seasonally adjust. Only the weighting was adjusted slightly to reduce the number of extreme values. Aviation freight is strongly influenced by trading-day effects. The moving holidays of Easter and Thanksgiving were also found to be significant factors. A very significant outlier was September 2001, which is to be expected after the 9/11 catastrophe. The value for September 2001 dropped 13% from the previous month in the seasonally adjusted series. Other extreme, or nearly extreme, outliers include: February 1993, June and September 2001, and September and October 2002. The last outlier in October 2002, 4.5 standard deviations from the mean, increased from the previous month by 25% after adjusting for seasonality. Aviation freight easily passed the diagnostic standards, M1 through M11. The value for months to cyclical dominance (MCD) was 3. A multiplicative model was used to model the Aviation Freight data series.

NOTE: For the February 2009 TSI, no data points were forecasted and used in the model.

Passenger. July and August are the highest months for passenger data. February is consistently the low month. The seasonality of aviation passenger traffic is very evident. The defaults in the seasonality adjustment software work well, except for the weighting. A slight adjustment was made in the latter as in the aviation freight data to reduce the number of extreme values. Both moving holiday and trading-day effects were significant in aviation passenger series.

Like aviation freight, aviation passenger was impacted strongly by 9/11, where there was a 30-percent drop after seasonality was adjusted. In fact, much of 2001 was unusual with outliers from July through October, and where the August through October outliers were extreme. The only other extreme outliers were back in February and March 1991. No other extreme outliers were uncovered. Aviation passenger data easily passed the diagnostic standards, M1 through M11. The value for MCD was 4, which is the median for the 10 series. A multiplicative model was used to model the aviation passenger data series.

Moving seasonality was detected at the one percent level. Although it was also detected in two other TSI components, natural gas and transit, it is never as significant or as consistent as in airline passenger.

NOTE: For the February 2009 TSI, no data points were forecasted and used for determining the seasonal factors.

[1] For a good overview of the X-11 method, see D. Ladiray and B. Quenneville. *Seasonal Adjustment with the X-11 Method*. 2001. New York: Springer-Verlag.Ashley, J. D. 2001. Why Seasonal Adjustment – Draft. Washington, D.C.: Bureau of the Census. Available online at http://www.catherinechhood.net/WhySeasAdj.pdf.

Findley, D. F., Monsell, B. C., Bell, W. R., Otto, M. C., and Chen, B. 1998. New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program. Washington, D.C.: Bureau of the Census. Available online at http://www.census.gov/ts/papers/jbes98.pdf.

Hood, C. 2009. Seasonal Adjustment and Time Series FAQ. Nashville, TN: Catherine Hood Consulting. Available online at http://www.catherinechhood.net/safaqmain.html.

Ladiray, D., and Quenneville, B. 2001. Seasonal Adjustment with the X-11 Method. New York: Springer.