RESEARCH REPORT SERIES
(Statistics #2014-09)

Effect of Trading Day Regressors on Seasonal Adjustment of Growth Rates

James Livsey
Osbert Pang
Tucker McElroy

Center for Statistical Research & Methodology
Research and Methodology Directorate
U.S. Census Bureau
Washington, D.C. 20233

Report Issued: September 5, 2014

Effect of Trading Day Regressors on Seasonal Adjustment of Growth Rates

James Livsey, Osbert Pang, and Tucker McElroy

Abstract

This report clarifies an issue involved in performing seasonal adjustment with growth rates. Specifically, we consider the effect of including or excluding trading day regressors in the seasonal adjustment routine. This work provides insight to X-12-ARIMA (or its successor X-13ARIMA-SEATS) users who wish to calculate a 12-month concatenated growth rate or any other derived function at some point before, during, or after seasonal adjustment. This may be a delicate issue since the operation of calculating an annual (or concatenated annual) growth rate can remove seasonality, but does not remove trading day effects.

1 Introduction

Using a 12-month concatenated (compounded) growth rate in tandem with seasonal adjustment is a practice at some central banks. For example, some published stock series at the Bank of England are produced from a compilation of internally defined adjustments, X-13ARIMA-SEATS (X-13A-S) seasonal adjustment, and 12 month concatenated growth rates.

This report originates from an inquiry to the U.S. Census Bureau from an employee at the Bank of England. The Bank of England uses a version of X-13A-S to perform seasonal adjustment. This report will assume some familiarity with the X-13A-S program; for more information, see [1]. In the inquiry, it was observed that using X-13A-S to seasonally adjust a level series led to discrepancies in the resulting 12-month concatenated growth rates. Specifically, these discrepancies appeared when comparing the growth rates of the non-seasonally adjusted (NSA) data or the seasonally adjusted (but not trading-day adjusted) data with those from the seasonally and trading-day adjusted data. Sample data with which these three growth rates could be calculated was provided. Figure 1 plots these different methods of obtaining growth rates. The methodology for constructing each of the lines in this graphic will be further discussed in Section 2.1. The goal of this report is to account for the discrepancies visible between the series plotted in red and the relatively similar series in black and blue.

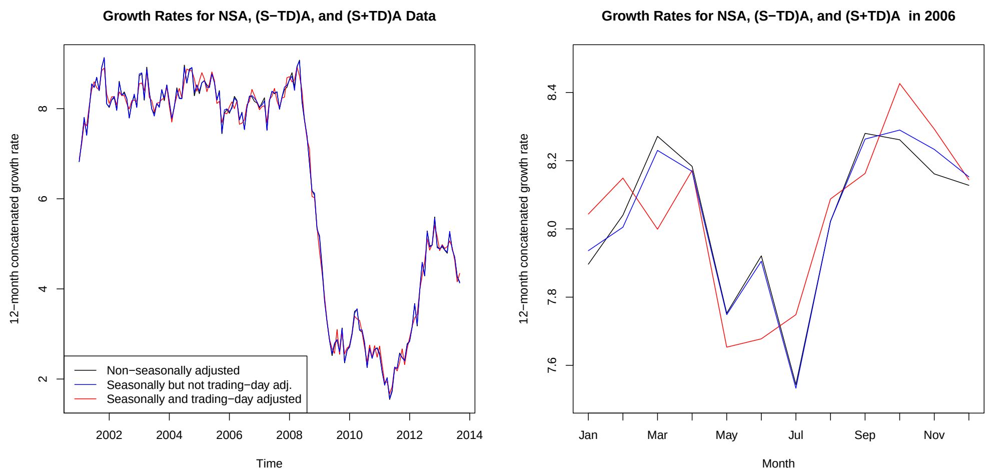

Figure 1: Plot of 12-month concatenated growth rates g_t^{12m} given by equation (3) for NSA and each seasonal adjustment method outlined in section 2.1, i.e., (S-TD)A and (S+TD)A.

The rest of the report proceeds as follows. In section 2, we describe the data set provided, outline the procedure used by the Bank of England during seasonal adjustment, and introduce notation to rigorously define the questions asked. In section 3, the solution is presented and verified in the spectral domain. Section 4 will contain concluding remarks.

2 Preliminaries

The data provided are non-seasonally adjusted (NSA) level and flow data for some monthly (stock) series. For confidentiality reasons, no other information about the data was provided. This data differs from classic level/flow data in that the first differences of the level series do not exactly equal the monthly flow. This discrepancy is reconciled via a recursively defined break index $\{b_t\}$ applied to both levels and seasonally adjusted flows. The outline of its use and definition are given in Figure 2. The technicalities of this break index, and its exact use, are further discussed in Appendix A. The calculations and results in this report do not hinge on the understanding of b_t.

Let $Y_1, Y_2, \ldots, Y_T$ represent the level data. The reported flow values (not the first difference of the level series) for month t, will be notated $flow_t$. The monthly levels and flows for the provided data are plotted in Figure 3.

Two other important derived quantities from the data are now discussed. First, let the fractional growth rate be generally defined as

$$f_t = \frac{flow_t}{Y_{t-1}}. \tag{1}$$

For the case of seasonally adjusted data, we can combine (1) and (†) from Figure 2, thus allowing

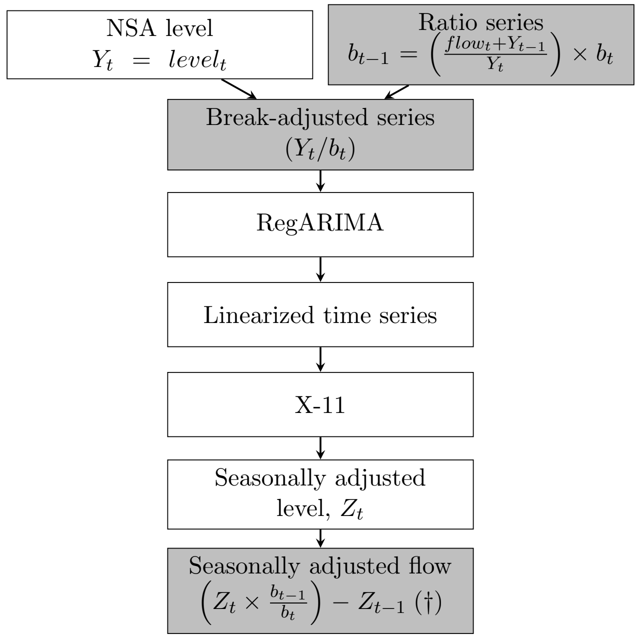

Figure 2: Bank of England's seasonal adjustment procedure, matching those found at `http://www.bankofengland.co.uk/`, but with notation adjusted for consistency; gray boxes only apply for levels series whose first differences do not equal flows (as in our situation).

us to express the fractional growth rate obtained from seasonally adjusted data as

$$f_t^{SA} = \frac{\left(Z_t \times \frac{b_{t-1}}{b_t}\right) - Z_{t-1}}{Z_{t-1}}. \tag{2}$$

Either way, the percentage change from month $t-1$ to month t is given by $100 \times f_t$. As an example, NSA fractional growth rates are plotted in Figure 4. We see that the fractional growth rates tend to be fairly small, ranging from a low of -0.5% to a high around 1.5%. Second, a 12-month concatenated growth rate is obtained from the 12 most recent fractional growth rates via

$$g_t^{12m} = 100\left[(1+f_t)(1+f_{t-1})\cdots(1+f_{t-11}) - 1\right]. \tag{3}$$

A more enlightening form of (3) may be

$$1 + \frac{g_t^{12m}}{100} = (1+f_t)(1+f_{t-1})\cdots(1+f_{t-11}). \tag{4}$$

This representation can be viewed as an 12-month rate of return or, in the case of non-negative f_t, as a compound interest analogue. Moreover, due to their small values, we may assume the products $f_t f_{t-j}$ are approximately zero for all t and j; thus, we have $1 + \frac{g_t^{12m}}{100} \approx 1 + f_t + f_{t-1} + \cdots + f_{t-11}$. The second expression is the seasonal period sum operator $U(B) = (1 + B + \cdots + B^{11})$, applied to f_t, which will strongly dampen and possibly eliminate seasonal movement. Hence, g_t^{12m} is not

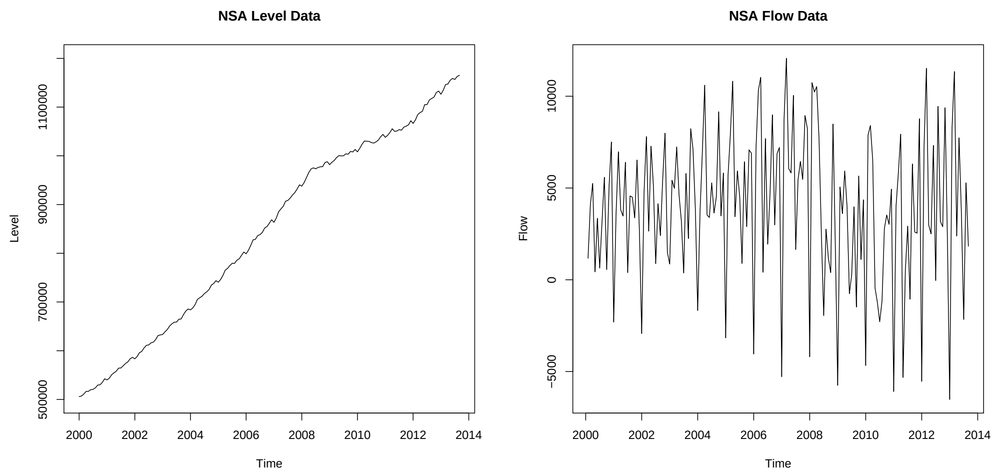

Figure 3: Monthly level and flow data for a sample series.

expected to be very seasonal, just as seasonally differenced series are not expected to display much seasonal behavior. This will hold whenever the data source for g_t^{12m} has small f_t We will present this line of reasoning more rigorously with a simple analogue in section 3.1.

2.1 Growth Rate Variants

We can now more carefully define procedures for producing the series displayed in Figure 1. First, let Z be used to denote seasonally adjusted levels data (e.g., the .d11 values returned by X-13A-S). If we allow "(S-TD)A" to refer to the case where a series is seasonally adjusted, but not trading-day adjusted, then we can denote the seasonally adjusted level series as $Z^{(S-TD)A}$ (to reflect the fact that trading-day movements may still reside within the seasonal adjustment). Similarly, let us refer to the case where a series is both seasonally and trading-day adjusted as "(S+TD)A," with the corresponding seasonally adjusted level series being $Z^{(S+TD)A}$. The three different series depicted in Figure 1 are calculated as follows:

1. NSA (black line) is produced by using NSA level and flow data to calculate fractional growth rates f_t^{NSA} via (1). These NSA fractional growth rates are then used to calculate the growth rate $g_t^{12m,NSA}$ using (3).

2. (S-TD)A (blue line) is just the seasonal adjustment (without a trading-day adjustment). The resulting seasonally adjusted series (the .d11 file output) are $\{Z_t^{(S-TD)A}\}$. The SA flow data is then calculated via (†) from Figure 2 using the $Z^{(S-TD)A}$ values. The corresponding fractional growth rates, and consequently the 12-month concatenated growth rate, are obtained via (1) and (3) using the seasonally adjusted levels and flows, and are labeled as $f_t^{(S-TD)A}$ and $g_t^{12m,(S-TD)A}$, respectively.

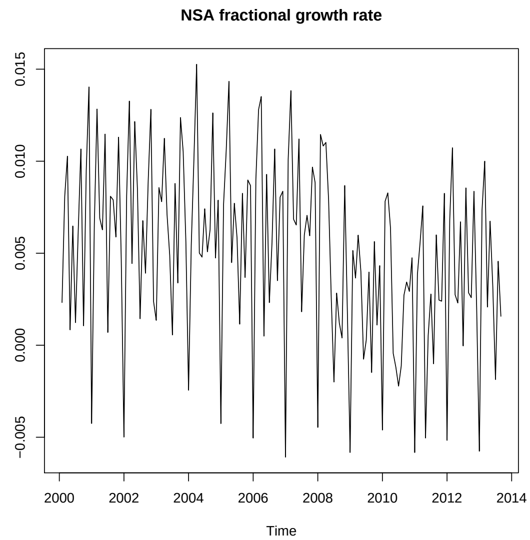

Figure 4: The NSA fractional growth rates

3. (S+TD)A (red line) differs from (S-TD)A only in the inclusion of a trading-day adjustment when performing the initial X-13A-S seasonal adjustment, thus yielding $\{Z_t^{(S+TD)A}\}$, $f_t^{(S+TD)A}$, and $g_t^{12m,(S+TD)A}$ as the seasonally adjusted levels and corresponding fractional and 12-month concatenated growth rates.

3 Answers to Inquiry

From the proximity of the black and blue lines in Figure 1, we might surmise that most seasonal movement is removed when taking a 12-month concatenated growth rate. Periodic movement that does not recur over a time span of 12 months, as would be the case for a trading day movement (with 28 year = 336 month periodicity), would still remain. Hence, a seasonal adjustment that accounts for a statistically significant trading day movement would result in a deviation from the NSA values when (3) is applied. In the same vein, if a trading day effect did exist in the original series, then a seasonal adjustment that failed to account for it would still display significant movement at the typical TD frequencies (for monthly series, 0.348 and/or 0.432 cycles per month). This could be the source of concern that prompted the initial inquiry.

3.1 A 12-month Change Analogue

Suppose the original series Y_t (logged if necessary) has trend T_t, seasonal S_t, trading day TD_t, and irregular I_t components that appear additively. In the case that TD_t does not appear prominently or at all, its estimated regression coefficients will be small or zero. For purposes of illustration, we suppose that S_t is deterministic, so it is completely annihilated by seasonal adjustment filters $(1-B^{12})S_t = 0$. This is not realistic, and our discussion could be adapted to the case where $U(B)S_t$ is stationary (with mean zero) instead of identically zero, but the fundamental analysis does not change. So, we may write

$$Y_t = T_t + S_t + TD_t + I_t, \tag{5}$$

which represents the NSA series; let us denote the NSA 12-month change as ∂Y_t:

$$\partial Y_t = (1 - B^{12})Y_t = (1 - B^{12})T_t + (1 - B^{12})I_t + (1 - B^{12})TD_t.$$

We might view this hypothetical ∂Y_t as an analogue for the 12-month concatenated rate $g_t^{12m,NSA}$ of the non-seasonally adjusted data. Let us also consider two perfect seasonal adjustments; the first explicitly models and removes the TD, while the second ignores the possible presence of TD. Denoting these seasonal adjustments by $Z_t^{(S+TD)A}$ and $Z_t^{(S-TD)A}$ respectively, we have

$$\begin{aligned} Z_t^{(S+TD)A} &= T_t + I_t \\ Z_t^{(S-TD)A} &= T_t + I_t + TD_t. \end{aligned}$$

Note that if TD is not actually present, then no discrepancy will exist between $Z_t^{(S-TD)A}$ and $Z_t^{(S+TD)A}$. This is an idealized scenario, however, as any seasonal adjustment would actually distort the trend and irregular components somewhat. Nevertheless, let us continue with this line of reasoning for didactic purposes. The corresponding 12-month changes are then

$$\begin{aligned} \partial Z_t^{(S+TD)A} &= (1 - B^{12})T_t + (1 - B^{12})I_t \\ \partial Z_t^{(S-TD)A} &= (1 - B^{12})T_t + (1 - B^{12})I_t + (1 - B^{12})TD_t, \end{aligned}$$

which we could view as theoretical counterparts to $g_t^{12m,(S+TD)A}$ and $g_t^{12m,(S-TD)A}$, respectively. The series as displayed in Figure 1 would conceivably appear quite similar to ∂Y_t, $\partial Z_t^{(S-TD)A}$, and $\partial Z_t^{(S+TD)A}$. In our idealized framework, there is no difference between the NSA 12-month change ∂Y_t and the SA-TD 12-month change $\partial Z_t^{(S-TD)A}$. The discrepancy between these series and the 12-month change $\partial Z_t^{(S+TD)A}$ is $(1 - B^{12})TD_t$, or ∂TD_t for short.

All that remains is to verify that ∂TD_t matches the discrepancies between $g_t^{12m,NSA}$ and $g_t^{12m,(S+TD)A}$ seen in data provided to us. Verifying this would effectively answer the inquiry. We perform a seasonal adjustment in X-13A-S, with trading day regressors included, on the NSA level series data. Since a log transformation is preferred, the resulting adjustment is multiplicative,

$$Y_t = T_t S_t TD_t I_t, \tag{6}$$

instead of the additive representation given in (5). However, a multiplicative representation is additive when viewed on a logarithmic scale, so this is a minor issue. The estimated trading

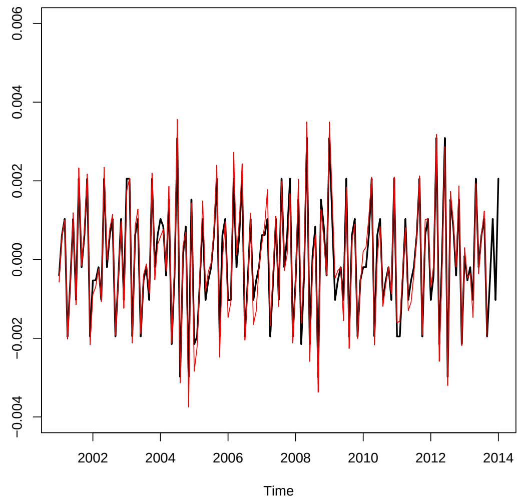

Figure 5: The 12 month difference of log of trading day effects (black) over-plotted with the discrepancies in 12 month growth rates between NSA and SA-No-TD (red). The sample correlation coefficient between the two series is 0.9829.

day effects $\{\widehat{TD}_t\}$ (the .d18 file from the X-13-A-S seasonal adjustment) are extracted from the X-13-A-S seasonal adjustment.

To mimic the 12-month growth rates, we use the seasonal difference operator $(1 - B^{12})$ and apply it to the log of the estimated trading day values:

$$(1 - B^{12}) \log(\widehat{TD}_t) = \log(\widehat{TD}_t) - \log(\widehat{TD}_{t-12}). \tag{7}$$

We then compare this series to the discrepancy formed by taking the difference between the 12-month concatenated growth rates for NSA and (S+TD)A. These two series are plotted in Figure 5, and it is clear these two series track each other nearly perfectly. Indeed, the correlation between the two series is 0.9829, which indicates that the main reason we observe a difference in the concatenated growth rates is the influence of the trading day effects in the NSA and SA-No-TD variants.

3.2 An examination of spectral densities

Section 3.1 used a seasonal difference operator to show that the discrepancies in the concatenated growth rates were largely a function of the trading-day effect that was not accounted for. We might consider examining spectral densities as a means of reaching the same conclusion. To start, (3) (or the equivalent expression (4)) gives the 12-month concatenated growth rate at a particular time as a function of the 12 most recent fractional growth rates, expressed as a percent. Taking logarithms,

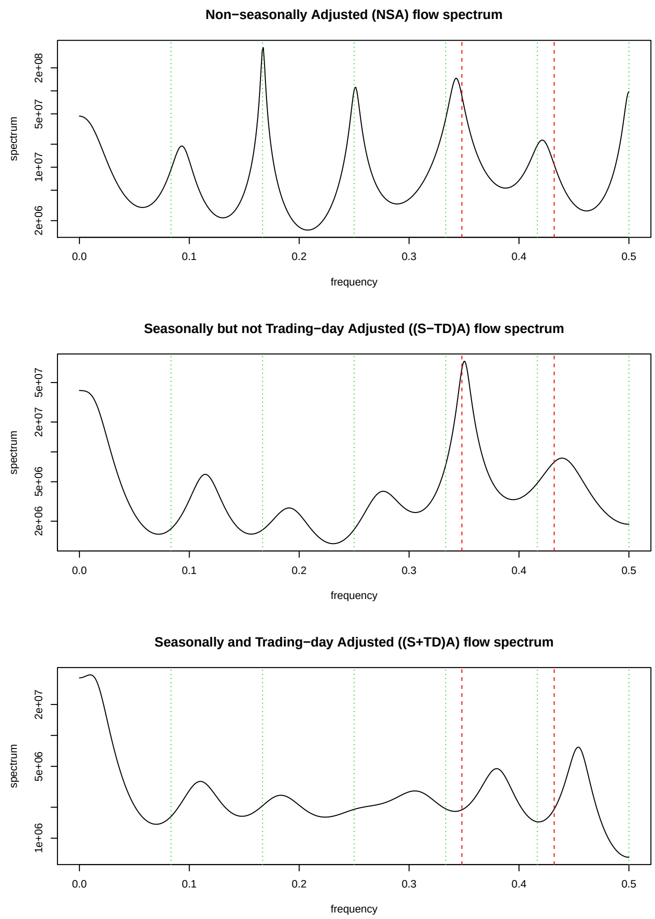

Figure 6: Estimated spectral densities of flow series: non-seasonally adjusted (top), seasonally but not trading-day adjusted (middle), and seasonally and trading-day adjusted (bottom).

this can be rewritten as

$$\log\left(1+\frac{g_t^{12m}}{100}\right)=\sum_{i=0}^{11}\log(1+f_{t-i})\approx\sum_{i=0}^{11}f_{t-i}, \tag{8}$$

where the right-hand-side approximation relies on the fractional growth rates being small. As we have seen in Figure 4, this assumption is tenable for this particular set of data. If this is not the case, we can ignore the approximation and consider just the equality portion of (8).

We begin by examining the spectral densities of the various flow series. In the NSA case, these flows are measured, but in both of the SA cases, their flows are derived from seasonally adjusted level data. Figure 6 shows the spectral densities for the NSA flow series (top), the (S-TD)A flow series (middle), and the (S+TD)A flow series (bottom). We place red dashed lines at the trading day frequencies (0.348 and 0.432 cycles per month), while green dotted lines are located at common seasonal frequencies. In the top and middle densities, we see that there is a peak around the 0.348

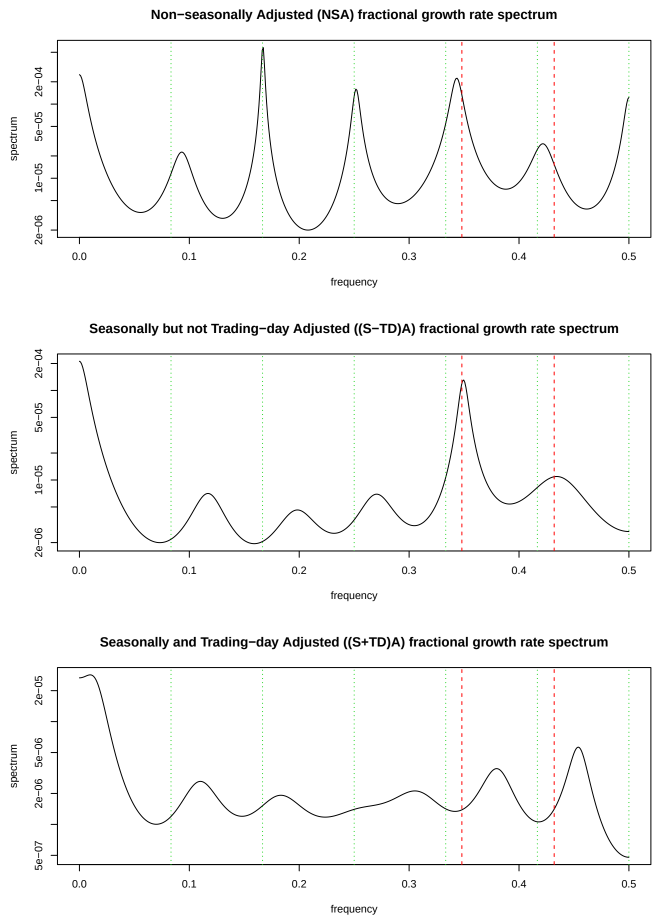

Figure 7: Estimated spectral densities of fractional growth rate series: non-seasonally adjusted f_t^{NSA} (top), seasonally but not trading-day adjusted $f_t^{(S-TD)A}$ (middle), and seasonally and trading-day adjusted $f_t^{(S+TD)A}$ (bottom).

trading day frequency, which is not present in the bottom density. This should be expected: the trading day effect is accounted for in the seasonal adjustment associated with the bottom density, but not for the one corresponding to the middle density.

The fractional growth rates defined in (1) are a function of flow and level data, so we might suspect that these same tendencies would appear in the spectral densities of the fractional growth rates for the three cases. Looking at Figure 7, our suspicions are indeed confirmed, as we again observe a peak in the spectral density at the 0.348 trading day frequency for the fractional growth rates produced using NSA (top) or (S-TD)A (middle) flows and levels, which is absent in the spectral density for (S+TD)A (bottom).

Lastly, if we examine Figure 8, we see that there is evidence of a trading day effect in the logarithm of the 12-month concatenated growth rate, but no strong indication of a discernible

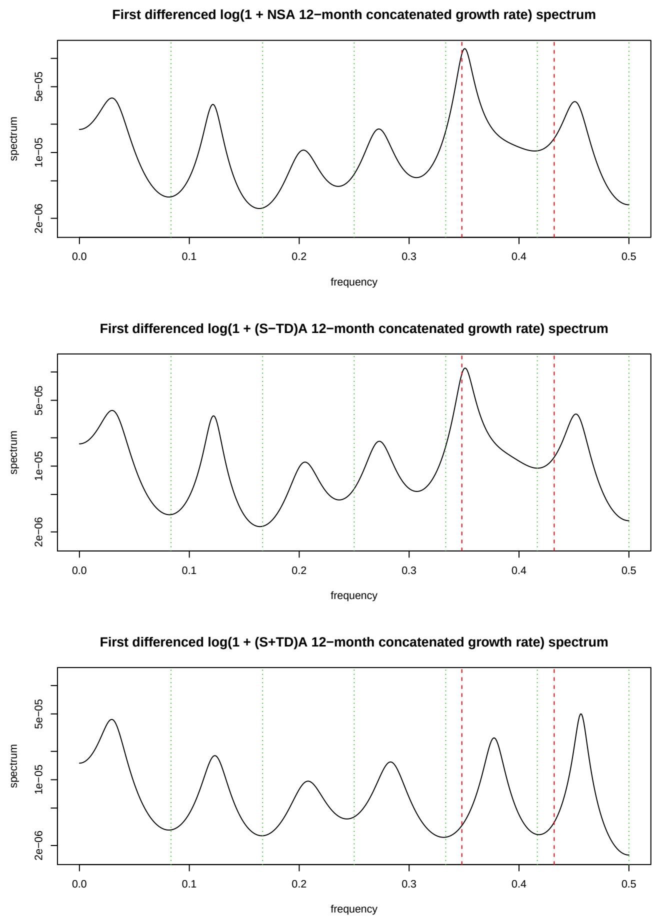

Figure 8: Estimated spectral densities of differenced log 12-month concatenated growth rates: non-seasonally adjusted $g_t^{12m,NSA}$ (top), seasonally but not trading-day adjusted $g_t^{12m,(S-TD)A}$ (middle), and seasonally and trading-day adjusted $g_t^{12m,(S+TD)A}$ (bottom).

seasonal pattern at the typical frequencies (since the right-hand-side approximation of (8) is the seasonal sum operator referenced earlier, the lack of a seasonal pattern should be expected). Given that the logarithmic function of g_t^{12m} shown in (8) is one-to-one, the spectral density of 12-month concatenated growth rates should possess similar features. Hence, it is apparent that the trading-day effect seen in Figure 8 exists in the 12-month concatenated growth rate calculated using the (S-TD)A levels and flows, but it will not be obvious in the 12 month concatenated growth rate calculating using the (S+TD)A levels and flows. Given that the month-to-month change in the log transform of the concatenated growth rate can be approximated as

$$\log\left(1+\frac{g_t^{12m}}{100}\right)-\log\left(1+\frac{g_{t-1}^{12m}}{100}\right)=\sum_{i=0}^{11}\log(1+f_{t-i})-\sum_{i=0}^{11}\log(1+f_{t-1-i})$$
$$=\log(1+f_t)-\log(1+f_{t-12})\approx f_t-f_{t-12}.$$

This means that this month-to-month change on a logarithmic scale is approximately the same as the year-to-year change in f_t for the corresponding calendar month. Hence, this operation eliminates any seasonal pattern that perfectly repeats over a 12-month period, but does not annihilate trading-day effects that exist in the data.

4 Conclusion

The various plots of the spectral densities reveal that the presence of a trading-day effect in the non-seasonally adjusted level and flow data carries over to the NSA fractional growth rates and consequently, the NSA 12-month concatenated growth rate. This might be broadly applicable to any series that is derived from NSA level and flow data. Incorporating trading day regressors into the seasonal adjustment removes the trading day effect from the SA-level data in this case, and since SA-flow data is a function of SA-level data, this has the effect of removing the trading day effect present in the NSA flow data. Thus, the trading day effect that exists in the NSA fractional growth rates induces a trading day effect in the NSA 12-month concatenated growth rate. These trading day effects, however, disappear for the fractional and concatenated growth rates produced using output from the seasonal adjustment that includes trading day regressors. These results corroborate the conclusion of the seasonal difference approach. In other words, any observable discrepancy between the concatenated growth rates produced using NSA data and those produced using SA with TD output is largely attributable to the fact that the trading day effects in the original data carry over to derived functions of that original data (and do not transfer when accounted for).

A Appendix

Instead of the classic level and flow relationship, a general formula for our data of the change in levels from month to month is given as

$$\text{Closing balance} \;=\; \text{opening balance} \;+\; \text{transactions} \;+\; \text{other changes in value (OCVA).} \tag{9}$$

The closing balance is the aforementioned level data at time t, while the opening balance is the corresponding level from time $t-1$. The measured flow discussed throughout the report is the transactions term in (9). To better understand the OCVA term, the Bank of England provides the following explanation:

> *Transactions means the net change in balances over a reporting period that is attributable to the economic or financial behavior of households, businesses or other entities in the sectoral definition – e.g., putting more money into a savings account, or repaying a business loan. Other changes in the value of assets (OCVA) means the effects of changes in definitions, statistical re-classifications, write-offs, revaluation effects and such like.*

The OCVA values from each month are generally not observed. As an aid in estimating the OCVA, the break index b_t, in month t, is defined recursively, starting with $b_T = 1$, using

$$b_{t-1} = \frac{Y_{t-1} + flow_t}{Y_t} \, b_t.$$

> *Transactions ("flows" or "changes") and growth rates data produced by the Bank of England generally reflect the application of adjustments in order to remove the impact of "other changes in value of assets (OCVA)," also known as "breaks." (See: Changes, flows, growth rates). Amounts outstanding ("levels") data, on the other hand, are usually not adjusted for breaks in this way.*

This expression can be rearranged to obtain the ratio of break index values

$$\frac{b_{t-1}}{b_t} = \frac{Y_{t-1} + flow_t}{Y_t}. \tag{10}$$

Thus, using (9) and (10), we get

$$\text{OCVA}_t = Y_t - (Y_{t-1} + flow_t) = Y_t - \frac{b_{t-1}}{b_t} = Y_t \left(1 - \frac{b_{t-1}}{b_t}\right), \tag{11}$$

but since the values of the break index are typically close to one, a reasonable approximation for OCVA would be

$$\text{OCVA}_t \approx Y_t \left(1 - \frac{1}{b_t}\right). \tag{12}$$

References

[1] Time Series Research Staff - Statistical Research Division (2013) X-13ARIMA-SEATS Reference Manual. `http://www.census.gov/srd/www/x13as/`.