Topics in time-series analysis

Models · Seasonal adjustment · Imputation

Compiled from session4.tex @ 2024-08-06 15:28:46+02:00.

Day 4: Seasonal decomposition diagnostics

Andreï V. Kostyrka
19th of April 2024

UNIVERSITÉ DU LUXEMBOURG

Presentation structure

1. Seasonal-adjustment quality evaluation

2. Visual diagnostics of seasonality

3. Seasonal adjustment in R

 © Andreï V. Kostyrka, Université du Luxembourg, 2024

Seasonal-adjustment quality evaluation

What should be diagnosed

Why is parametric RegARIMA important if X13 is iterated ‘non-parametric smoothing’ = weighted averaging?

- A model is used to eliminate *(pre-adjust)* the deterministic calendar component
- A model is used to deal with influential observations (in the pre-treatment, remove aberrant effects long before the down-weighting in the X11 smoothing starts)
- A model is used to extend the ends of samples to avoid biases or strong revisions

RegARIMA diagnostics

Only residual diagnostics: X13 assumes a Gaussian ARIMA.

- Test residual properties: asymmetry and heavy tails (estimate the skewness and kurtosis), autocorrelation (2 lags or 2 seasonal lags = Ljung–Box)
- Failure = nothing can be done; affects calendar pre-adjustment and extrapolations at the ends

The final model is (semi-)auto-selected; more diagnostics should be carried out at a further step (after the non-parametric part).

Diagnostic M statistics

- Every X13 run is accompanied by a set of M quality statistics and an aggregate Q statistic (1978)
- Each statistic represents badness (the higher, the worse); the common acceptance threshold is $M_i < 1$
- Some statistics are misleading, **most can be safely ignored**
- Q and Q_{-M_2} are weighted averages of M_i and can be misleading, too
    - Ignore Q in favour of Q_{-M_2}

Meaning of statistics

- M_7: is there stable and identifiable seasonality? If not, no adjustment should be carried out
  - Adjustment can be enforced if the problem is stability, not existence of seasonality
  - Trust the plot more (M_7 is unreliable)
- M_8–M_{11} show how stable the seasonal patterns are
  - Better alternative: look at the plot
- M_1 show how noisy the irregular is (accuracy)
- M_2, M_3, M_5 are trend-related (ignore)
  - Historical legacy for RBC analysts
- M_4 is autocorrelation-related (ignore)
- M_6 is related to the smoothing window (ignore)

Joint significance testing

$$Y_t = \alpha_0 + X_t'\beta_0 + Z_t'\gamma_0 + U_t, \quad \mathbb{E}(U \mid X, Z) = 0$$

Consider $\mathcal{H}_0 \colon \gamma_0 = \gamma^*$. Under the true null, the **Wald statistic**

$$\hat{W} \coloneqq (\hat{\gamma} - \gamma^*)'(\widehat{\text{Var}}_{\text{HAC}}\, \hat{\gamma})^{-1}(\hat{\gamma} - \gamma^*) \overset{T\to\infty}{\sim} \chi^2_{\dim \gamma^*}$$

Reason: $\hat{\gamma}$, being a MM estimator, is asymptotically normal.

$\hat{W}$ is related to the popular F statistic: in finite samples, $\hat{F} \coloneqq \hat{W} / \dim \gamma_0 \sim \mathrm{F}_{\dim \gamma_0, n}$ is slightly more conservative.

NB: like $\hat{t} \coloneqq \hat{\beta} / \widehat{\text{SD}}(\hat{\beta})$ is **not** exactly Student-t-distributed, $\hat{F}$ does **not** follow the exact Fisher distribution law!

It is completely safe to use $k \cdot \mathrm{F}_{k,n} \overset{T\to\infty}{\sim} \chi^2_k$ (the accuracy of this approximation is the tamest of our problems).

Historical confusion about F statistics

Many textbooks define $\hat{F}$ via sums of residuals, R^2 etc. These primitive variants of $\hat{F}$ require conditionally homoskedastic Gaussian WN errors (too strong!).

- ANOVA is a relic of the past – forget it
  - Equivalent regression-based linear Wald tests work with heteroskedastic ergodic processes
- Two-sample *t*-tests and RSS-based *F*-tests are very restrictive and **generally invalid** versions of the less assuming Wald test (especially in TS context)
  - Use any consistent $\widehat{\mathrm{Var}}\, \hat{\gamma}$ to get valid results – impossible in sum-of-squares-based versions
- Unfortunately, the default output in regression summaries is non-robust – **do something** about it

Regression-based hypothesis testing

- Formulate hypotheses as properties / constraints in linear models and test them via the Wald statistic
- Test the zero mean of a stationary time series via the linear model $Y_t = \mu + U_t$
  - $\mathcal{H}_0\colon \mu = 0 \Rightarrow$ regress Y on the constant, get $\widehat{\mathrm{Var}}_{\mathrm{HAC}}$ (available in most packages), construct the t statistic
- Test the difference in levels of two time series, X_t and Y_t, via $Z_t \coloneqq X_t - Y_t$ and $\mathcal{H}_0\colon \mathbb{E}Z_t = 0$
- Vuong's test for non-nested models: compare the equality of two models' goodnesses of fit by regressing the difference of likelihood series on a constant
  - Calvet & Fisher (2004) suggest HAC VCOV estimation
  - Compare the equality of AICs, BICs etc. by adding a penalty to the likelihood series

Stable and moving seasonality tests

Define $SI_t := S_t + I_t$ or $S_t \cdot I_t$ (seasonal + irregular = original − trend − calendar). Let p index the period (e. g. 1,…,12, 1,…,12, … for monthly data).

1. $SI_{yp} = \alpha + D_p'\beta + U_{yp}$
   Test $\beta = 0$ yields F_S (significance of seasonal dummies)
   High F_S = there is seasonality
2. $SI_{yp} = \alpha + D_p'\beta + D_y'\gamma + U_{yp}$
   Test $\gamma = 0$ yields F_M (significance of annual dummies)
   High F_M = seasonality is unstable

X13 uses the ranks of SI_{yp} (Kruskal–Wallis, Friedman).

Computation of M_7

$$M_7 := \sqrt{\frac{1.5F_M + 3.5}{F_S}}$$

- M statistics say **nothing** about calendar effects
  - The series come already pre-adjusted if TD effects were found according to AICc (no C_t in SI_t)
- $M_7 > 1$ does not say why the seasonality is not identifiable (not stable or none at all?)
- M_7 is unreliable if I_t is not white noise!
- $M_7 < 1$ can be good, $M_7 > 1$ can be bad, but both are inconclusive (not a rigorous statistical test)
- Risk of skipping the necessary SA based on M and Q

Example: Belgian SP95 prices (1/4)

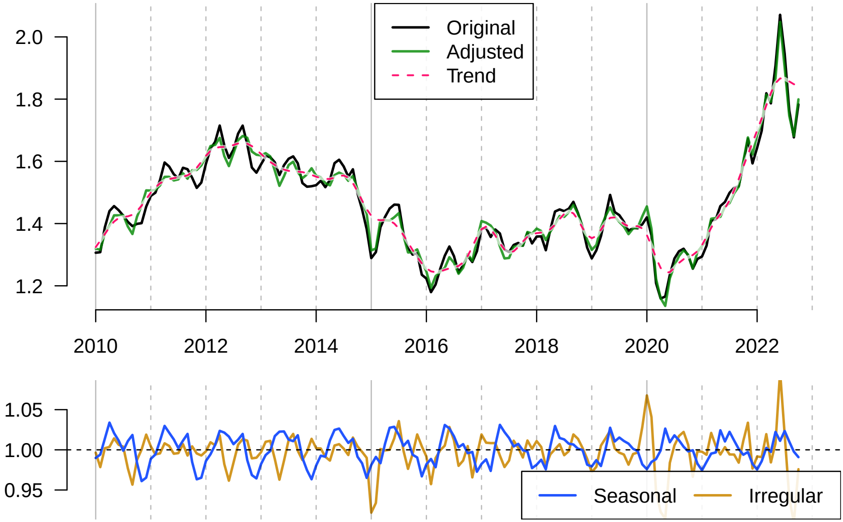

Example: Belgian SP95 prices (2/4)

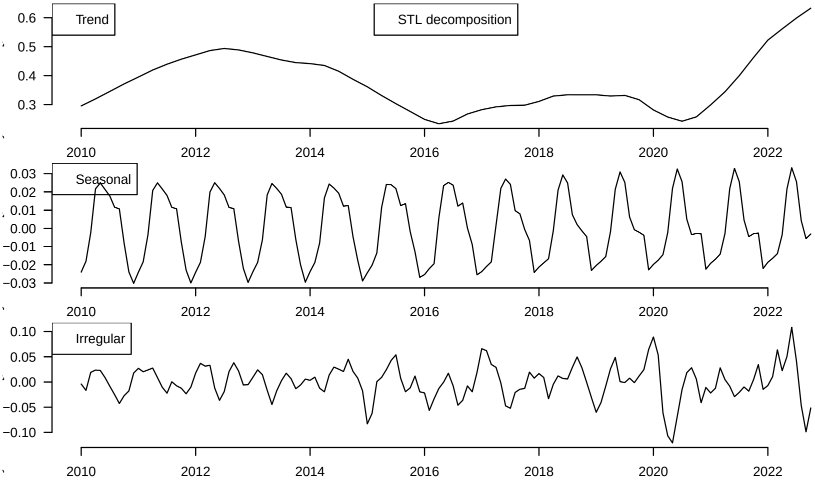

Example: Belgian SP95 prices (3/4)

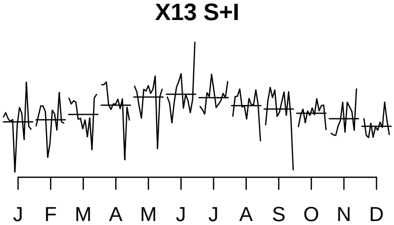

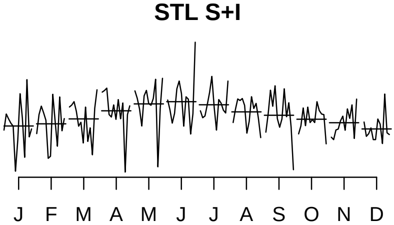

Regression-based (HAC-robust) joint significance tests for 11 period dummies (F_S) and annual dummies (F_M):

- X13 output: $F_S = 13.2$ ($p < 10^{-16}$), $F_M = 0.87$ ($p = 0.58$)
- STL output: $F_S = 2.9$ ($p = 0.0015$), $F_M = 1.05$ ($p = 0.41$)

Based on these plots and statistics, would you conclude that there is seasonality in petrol prices in Belgium?

Example: Belgian SP95 prices (4/4)

Monitoring and Quality Assessment Statistics

M-1	1.963	The relative contribution of the irregular over three months span
M-2	0.176	The relative contribution of the irregular component to the stationary portion of
M-3	0.135	The amount of period to period change in the irregular component as compared
M-4	2.263	The amount of autocorrelation in the irregular as described by the average durati
M-5	0.478	The number of periods it takes the change in the trend to surpass the amount of
M-6	0.571	The amount of year to year change in the irregular as compared to the amount o
M-7	1.253	The amount of moving seasonality present relative to the amount of stable seaso
M-8	1.489	The size of the fluctuations in the seasonal component throughout the whole ser
M-9	0.726	The average linear movement in the seasonal component throughout the whole
M-10	1.370	The size of the fluctuations in the seasonal component in the recent years
M-11	1.182	The average linear movement in the seasonal component in the recent years

M-Statistics

q: Bad (1.003)

q2: Bad (1.105)

summary

Good

Basic checks

definition: Good (0.000)

annual totals: Good (0.001)

Regarima residuals

normality: Uncertain (0.099)

independence: Good (0.845)

spectral td peaks: Good (0.546)

spectral seas peaks: Good (0.434)

Outliers

number of outliers: Good (0.000)

Residual seasonality tests

Qs test on SA: Good (1.000)

F-Test on SA (seasonal dummies): Good (0.922)

Qs test on I: Good (1.000)

F-Test on I (seasonal dummies): Good (0.860)

Misleading X13 output in lab conditions

Suppose that in years 2001–2010, monthly $\{Y_t\}$ were generated by the following SARMA DGP:

$$(1 - 0.3L)(1 - 0.4L^{12})Y_t = (1 + 0.5L)(1 + 0.95L^{12})\varepsilon_t - 1 \cdot \#_t^{\text{Mond}} + 1 \cdot \#_t^{\text{Frid}} + 5 \cdot \mathbb{I}(\text{Month}_t = \text{Dec}), \quad \varepsilon_t \sim t_3$$

Strong Christmas effect, strong week-day effect, perfectly stationary SARMA with finite variance…

SD(cal. + seas. effects) ≈ 26% of SD(Y_t) – should be adjustable, right?

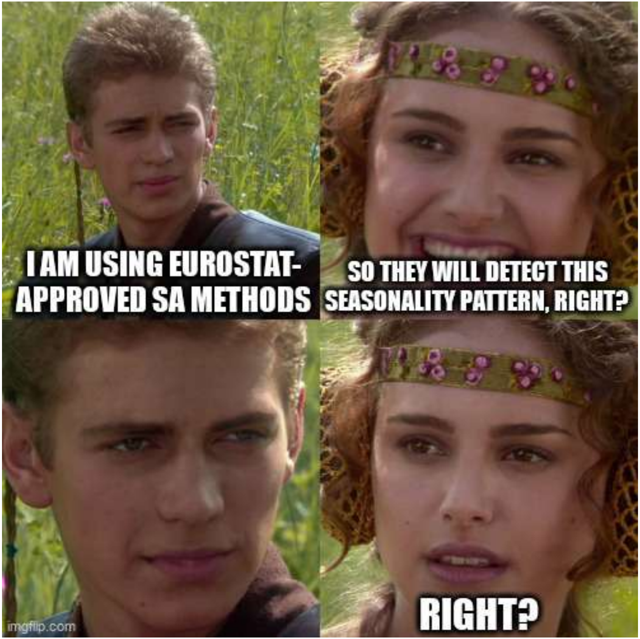
I AM USING EUROSTAT-
APPROVED SA METHODS
SO THEY WILL DETECT THIS
SEASONALITY PATTERN, RIGHT?
imgflip.com
RIGHT?

Synthetic seasonal series visualised

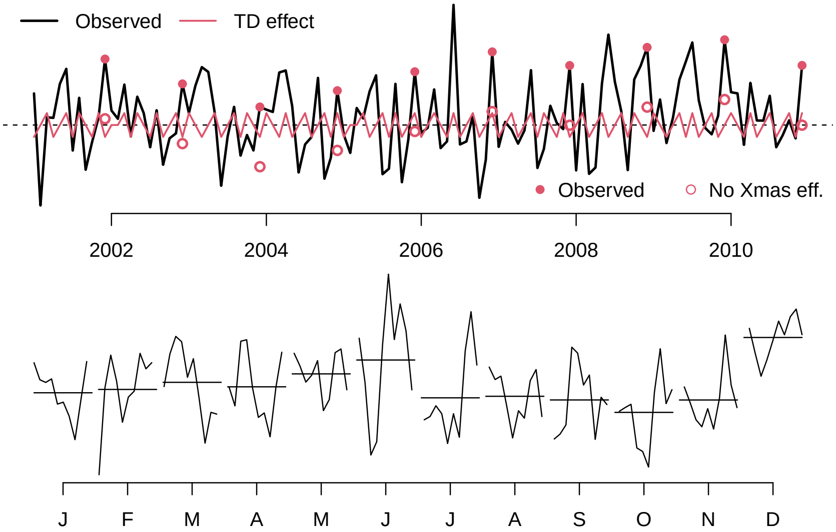

Synthetic seasonal series analysed

$M_7 = 1.21$, $Q_{-M_2} = 1.54$ (i. e. 'no seasonality, the adjustment quality is bad'). Why?

Despite the fact that…

- The December levels are above other months
- The auto-guessed seasonal model is correct
- The TD effect estimates are accurate

…the quality statistics in JDemetra+:

- Are uninformative about week-day effects
- Rely on non-robust VCOV (assuming homoskedastic + uncorrelated innovations): no HAC existed in 1975

Problem: distortion of M_7

$$M_7 := \sqrt{\frac{1.5F_M + 3.5}{F_S}}$$

	F_S	F_M	M_7
X13 rank ANOVA (no HAC)	4.7	0.23	0.91
Rank ANOVA + HAC	34.0	0.28	0.46
Huber WLS + HAC	16.7	0.24	0.48
OLS + HAC	17.1	0.36	0.49

M_7 can be **inflated** (if $\mathrm{Cor}(I_t, I_{t-1}) \gg 0$) due to a greater under-estimation of stable seasonality. Likewise, $\mathrm{Cor}(I_t, I_{t-1}) \ll 0 \;\Rightarrow\; M_7$ unrealistically small.

Why ignore most M statistics?

- Many M statistics test a strong non-constant trend ⇒ mean-stationary processes are marked as 'too noisy'
  - Only relevant for trend-cycle extraction
- The joint significance test (F statistic) assumes uncorrelated homoskedastic errors
  - But I_t should be well-approximated by an ARMA process with arbitrary autocorrelation pattern
  - Think critically and act from the first principles (transparency and general logic ≫ pre-baked numbers from 50-year-old algorithms)

JD+ sources: `Mstatistics.java`, `SeasonalityTest.java` (proof that the joint tests are **not** robust to autocorrelation even in the latest version!)

Output series code names

X13 provides several output tables. Plot the series together to ensure decomposition adequacy.

Name	Meaning
D10	seasonal component S_t
D11	seasonally adjusted series $Y_t - S_t - C_t = T_t + I_t + P_t$
D12	trend T_t
D13	irregular I_t
D16	seasonal + calendar $S_t + C_t$
D18	calendar C_t

Other numerical indicators

Q_S (Ljung–Box test for 2 seasonal autocorrelations): low p-value = evidence in favour of seasonality

- The original series should have $p < 0.01$
- The adjusted and irregular series should have $p > 0.10$
- Check pre-adjusted original (without outliers = EV-adjusted), EV-adjusted SA, and EV-adjusted irregular

Do not use the Box–Pierce test (worse inference).

Visual diagnostics of seasonality

Spectrum analysers

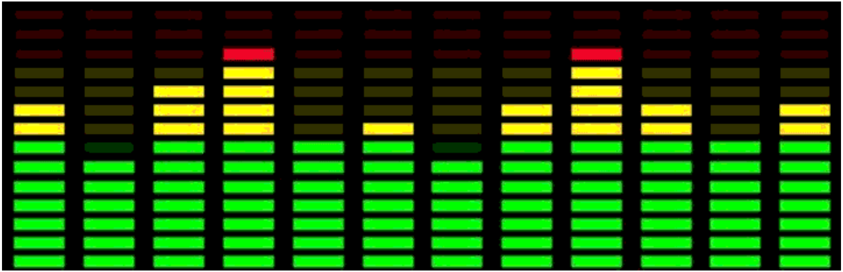

Spectrum analysers

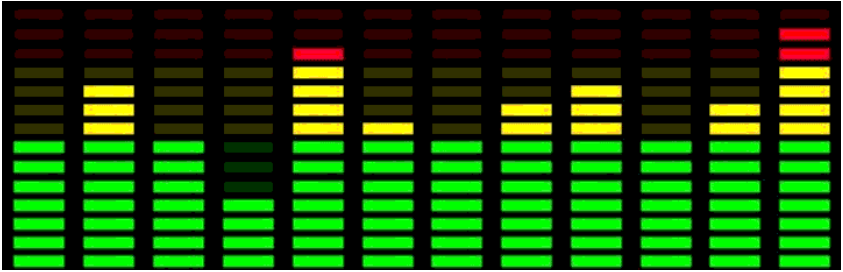

Spectrum analysers

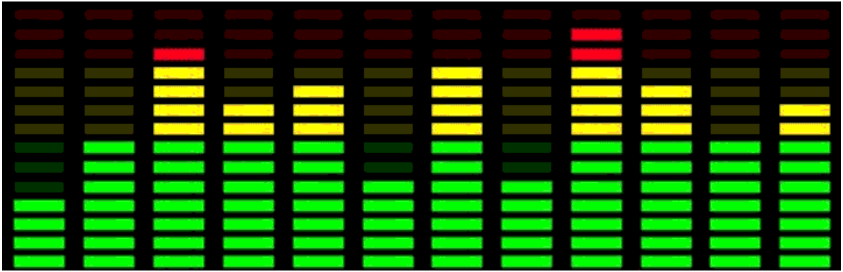

Spectrum analysers

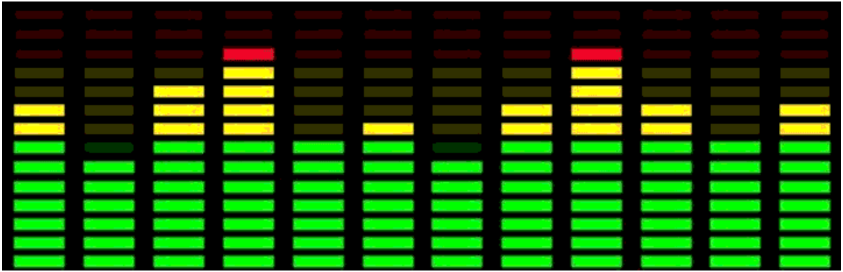

Spectrum analysers

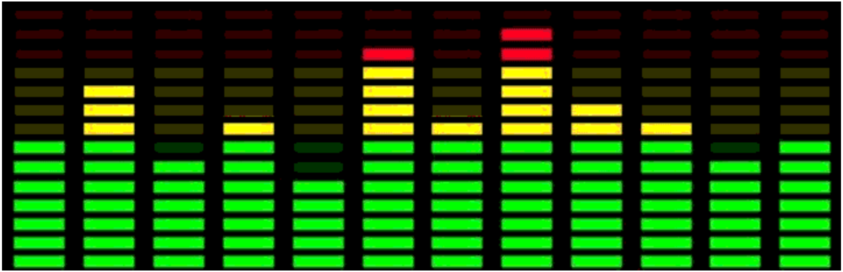

Spectral diagnostics

If $\{Y_t\}$ has any deterministic features, then a visualisation of the frequencies of $\{Y_t\}$ as should have spikes at known points.

- Spectral plots are easy to read (just like the audio spectrum analyser)
  - If a musical piece in A minor, then 440-Hz (or 2^n-multiples thereof, $n = 1, 2, 3, \ldots$) spikes are prominent
  - Everybody used a bass booster in their car / stereo system: same principle
- Shows **both** seasonal fluctuations and trading days
- Inference: visual significance (no numbers)

Fourier decomposition

Idea: convert the time series (time domain) into a set of frequencies and amplitudes of a set of sine waves that, when added together, reconstruct the original signal.

These videos explain the Fourier transform better than any textbooks:

- Visual introduction into Fourier transform
- Fast Fourier Transform

Spectrum estimation via DFT in R

A series with T observations $\{Y_t\}_{t=1}^T$ is converted into a set of T frequencies $\{f_t\}_{t=1}^T$ returned by the discrete transform:

$$f_k \coloneqq \frac{1}{T} \sum_{t=1}^{T} Y_t \cdot \exp\left(2\pi i \frac{t-1}{T} \cdot k\right)$$

Spectral density: $|f_k|^2$, i. e. the intensity of all frequencies.

For $\{Y_t\}_{t=1}^T$, `stats::fft(Y)` computes complex $\{f_k\}_{k=1}^T$.

Since $\exp(ix) = \cos x + i \sin x$, the spectral density (the intensity of all sine waves that, when added up, yield the original series) is `abs(stats::fft(Y))^2`.

How to read spectra

'What if I played the time series as a sound wave?'

Check the spikes at a pre-defined set of horizontal values.

- Spikes at 1/12, ..., 6/12 indicate phenomena occurring 1, ..., 6 times per year
  - 1/4 and 1/2 for quarterly data
- Spikes at trading-day frequencies = weekly effects
- Default in JDemetra+ (always plotted first)

Statistical tests for peak significance exists, but are obscure, non-standard, and do not tell more than visual analysis ('is the value at $k/12$ higher than the adjacent values').

Example: spectrum of a guitar string

Audio spectrum analysis in the open-source Audacity tool.

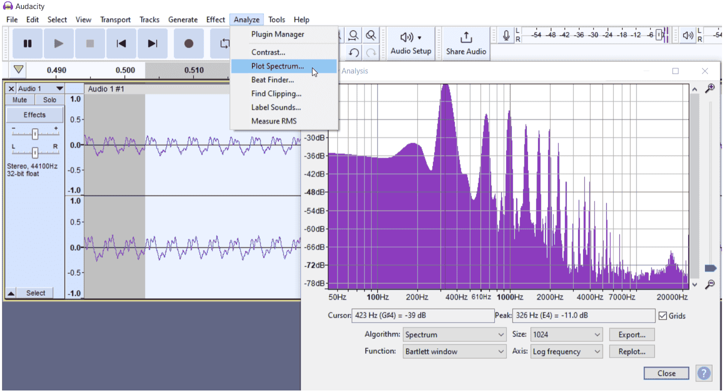

Note the harmonics at 660 Hz, 1320 Hz etc.: the string has stationary nodes due to real-world physical phenomena.

Spectrum and auto-correlations

So far, there have been very few references to ACFs, PACFs, their plots etc. Reason:

Power spectrum ≡ Fourier transform of the ACF.

It is much easier to diagnose time series (including non-stationary ones) with spectra than with ACF (which makes sense only for 2^{nd}-order-stationary ones).

- Low-frequency spikes ⇒ strong trend (long waves)
- Typically, one wants to detect both seasonal and calendar effects, and spectral analysis allows the researcher to see the presence of both in one picture

Spectrum of trigonometric monthly effects

If seasonality is due to monthly effects that are represented by a perfect sine wave, then expect spikes at the frequency 1/12 (‘a sine wave once per year’).

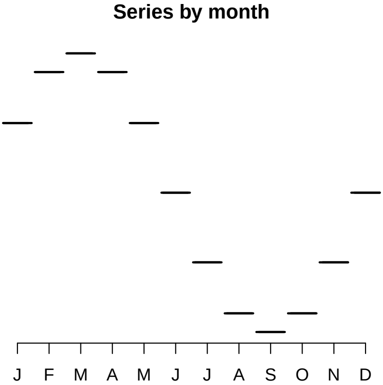

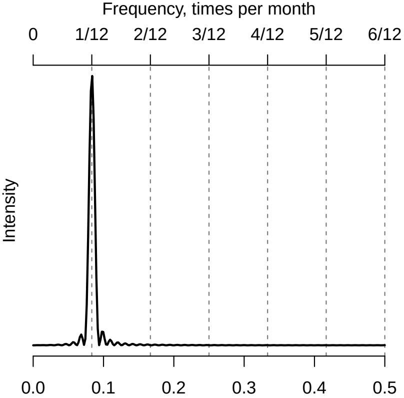

Spectrum of a one-month effect

If seasonality is due to one monthly effect, then expect spikes at **some** frequencies 1/12, ..., 6/12.

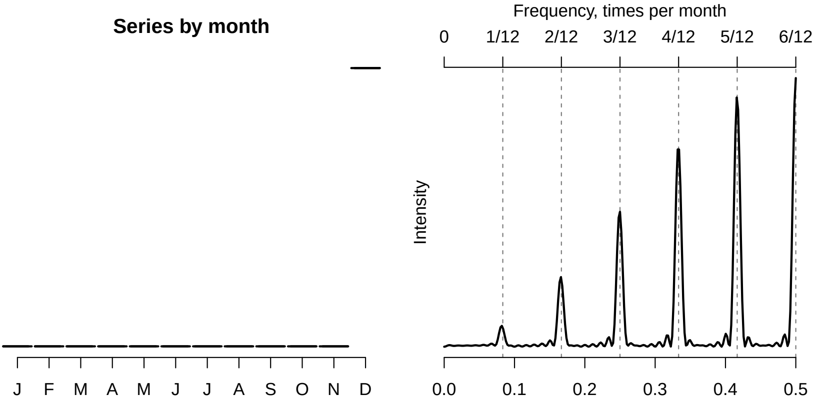

Spectrum of many monthly effects

If seasonality is due to varying monthly levels (12 different values), expect spikes at **some** frequencies 1/12, ..., 6/12.

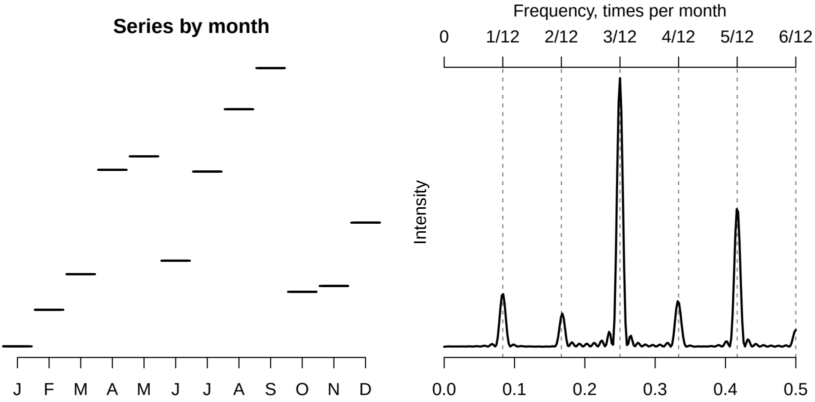

Nyquist–Shannon theorem and aliasing

If Y_t consists of sine waves with frequencies at most k, computing the exact spectrum of Y_t requires $T \geq 2k$ points.

With T points, it is possible to compute the weights of sine waves with periods ≥ 2 that add up to the original signal.

Example: A cycle that occurs 16 times per year cannot be detected with monthly data, but a 3-month cycle (4 times per year) can. Sampled 12 times per year, these two *look* identical (the high-frequency component is **aliased**):

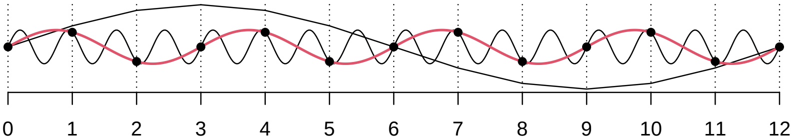

Weekly effects in monthly spectra

Weekly cycle effects should be undetectable with monthly data ($365.25/7 \approx 52.18$ cycles per year > 6). However...

- Weekly 'waves' are not in sync with monthly 'waves'
- 365.25 / 12 = 30.44 days/month = 4**.348** weeks/month
- If there is a weekly pattern, then every month, it is shifted by 0.348 weeks
  - 0.348 cycles / month = 4.179 cycles / year = detectable
  - The weekly wave is **aliased** by a longer wave that does not coincide with 2-month, ..., 12-month cycles ⇒ expect a new spike in the spectral diagram
- Other calendar frequencies are related to the 336-month calendar cycle and are aliased by several other long-waves – see Cleveland & Devlin (1980)

Spectrum of trigonometric weekly effects

If seasonality is due to smooth weekly fluctuations (i. e. effects of Monday, …, Sunday form a perfect sine wave), expect spikes at 0.348 (strong) and 0.432 (weaker).

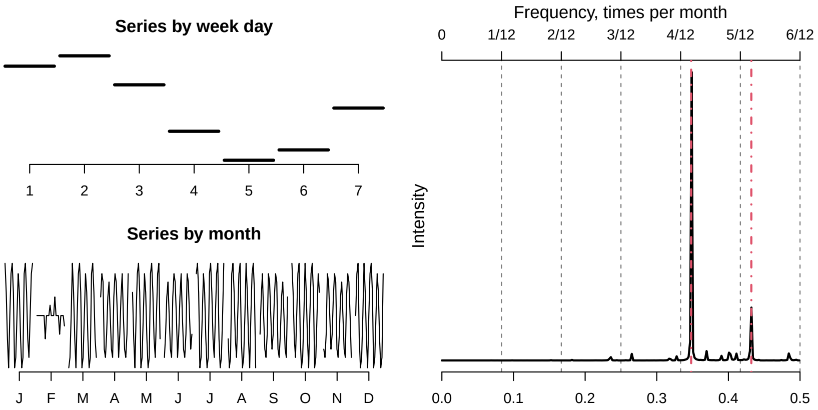

Spectrum of 1 week-day effect

If seasonality is due to the (A, A, A, A, A, B, B) week, then the spikes are the same.

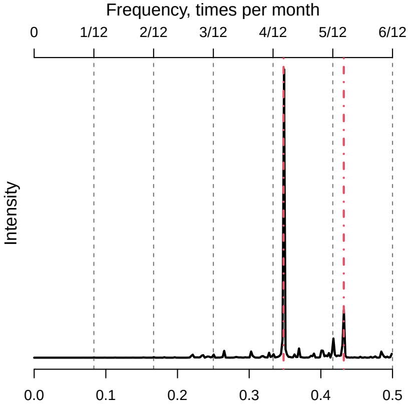

Spectrum of 7 week-day effects

If seasonality is due to the $(A_1, A_2, A_3, A_4, A_5, A_6, A_7)$ week, then the spikes are the same.

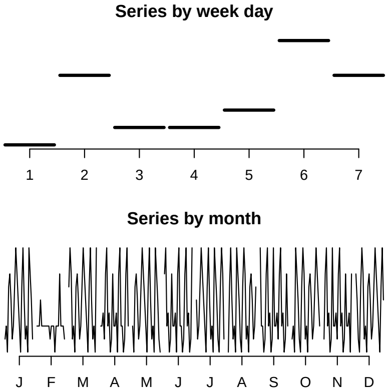

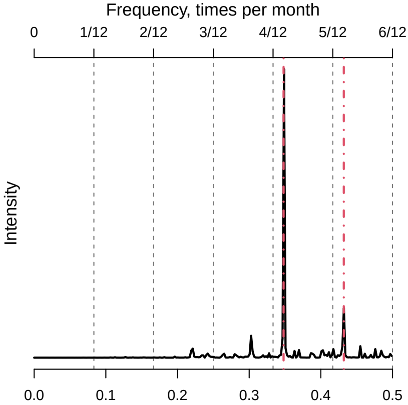

Spectrum of semi-weekly effects

If seasonality is due to the (A, B, B, A, B, B, B) week (or similar), then expect 2 more spikes: 0.304, 0.220.

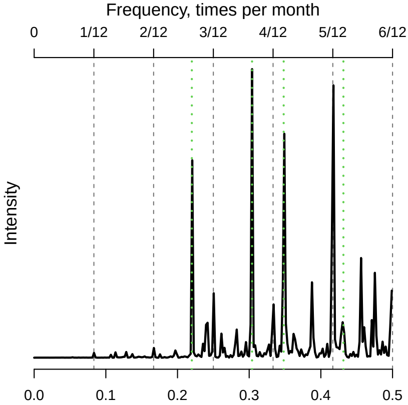

Rule of thumb for spectra

- Adjust the series if the seasonal / trading-day peaks of the spectrum are higher than the tallest non-seasonal and non-TD peak
- If the seasonal / TD peaks are not the strongest ones, do not adjust
- If the seasonal / TD peaks remain in the 'adjusted' series, check the RegARIMA, outliers etc.

Bad case: spectra (top: original, bottom: SA)

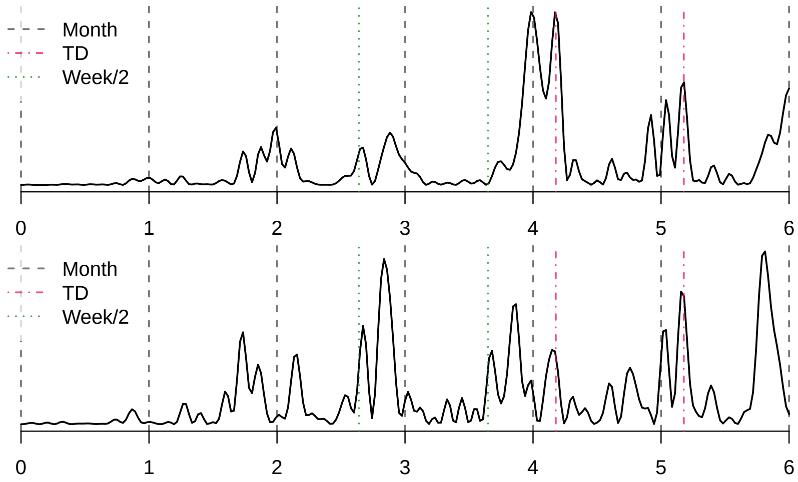

(click here to see the original series)

Very strong peaks may mask other effects

Example: airline passenger data (1976). Left: original, right: SA (original – $\hat{S}_t$) (no RegARIMA pre-treatment via TD). Note the calendar peaks in the right spectrum.

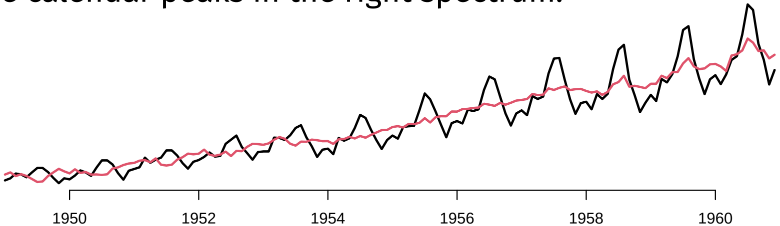

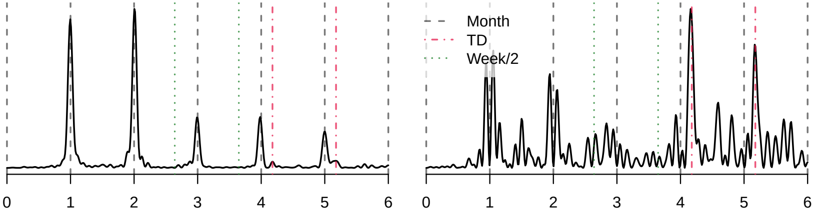

Working with quarterly data

Should anything change in the suggested procedure with quarterly data?

- No, same methods and parameters
- Relative variation in 90–92-day aggregates weaker than in monthly ones (28–31) + shorter series = expect more parsimonious models

Spectra of quarterly data

If seasonality is due to quarterly effects, then expect spikes at 1/4 or 1/2.

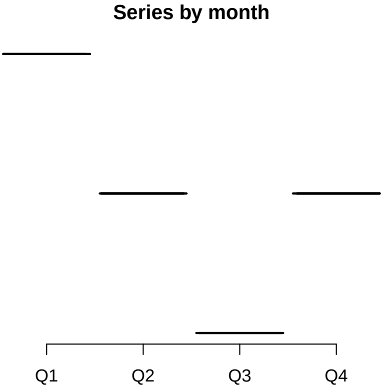

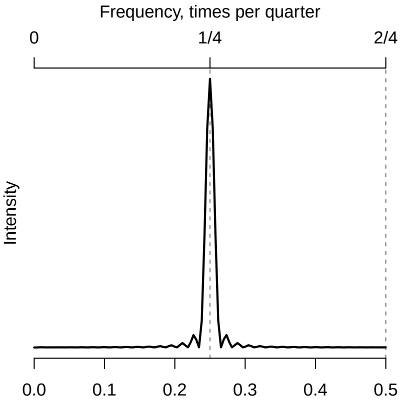

Spectrum of a one-quarter effect

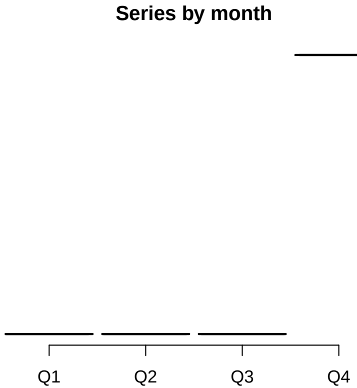

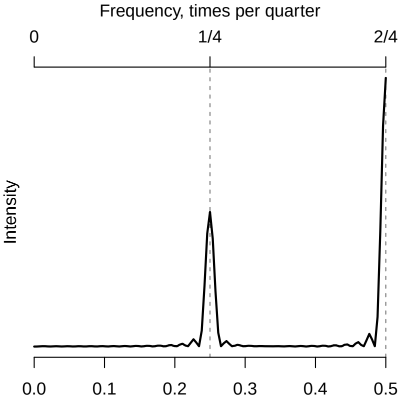

Spectrum of TD effects in quarterly data

TD effects create spikes at some of the 5 known frequencies; Cleveland & Devlin (1980) provide the exact values.

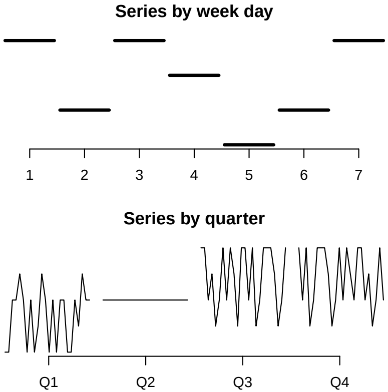

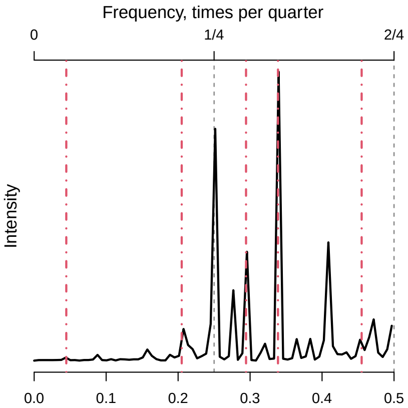

Time series length

No hard and fast rules as to how many years of data are necessary to obtain a publishable adjustment.

- <5 years: very short; should be used with strong warnings and only if the series are crucial to adjust
- 5–10 years of data: acceptable; when more data are added, there may be revisions, changes of model, recognition of previously undetected breaks
- 10 years: good
- >10 years: may have discontinuities or large changes in seasonal patterns; better split at breaks

Sliding spans for long series

- Choose a reasonable length limit (10–20 years)
- Split the sample into overlapping chunks
- Blend the adjusted series together
- Summarise the differences in overlaps

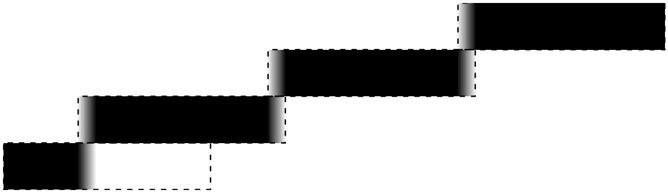

Analysing stability in long series

Although the sliding-span diagnostics are quite common, Eurostat recommends a simpler set of revision statistics.

The most useful ones:

- Bias: mean difference
- Dispersion: mean abs. difference, mean abs. percentage difference, standard deviation of the difference
- Proportion: $\frac{\text{Mean abs. diff.}}{\text{Mean abs. SA series}}$, $\frac{\text{SD(diff.)}}{\text{SD(SA series)}}$

Large discrepancies = unstable SA = less reliable analysis.

Choosing SA rules in long series

Imagine a scenario: in different subsamples, the seasonality strength indicator is different, the leap year, auto-selected model etc. is different.

Period	M_7	Trad. day	Leap year	ARIMA
1998.01–2007.12	1.1	1 WD	No	1,0,0
2004.01–2013.12	0.5	–	Yes	1,0,0
2013.01–2022.12	0.7	–	Yes	0,0,1

Should one adjust these series with the auto-selected rules ($n = 300$, $n_{\max} = 120$)?

It was a trick question

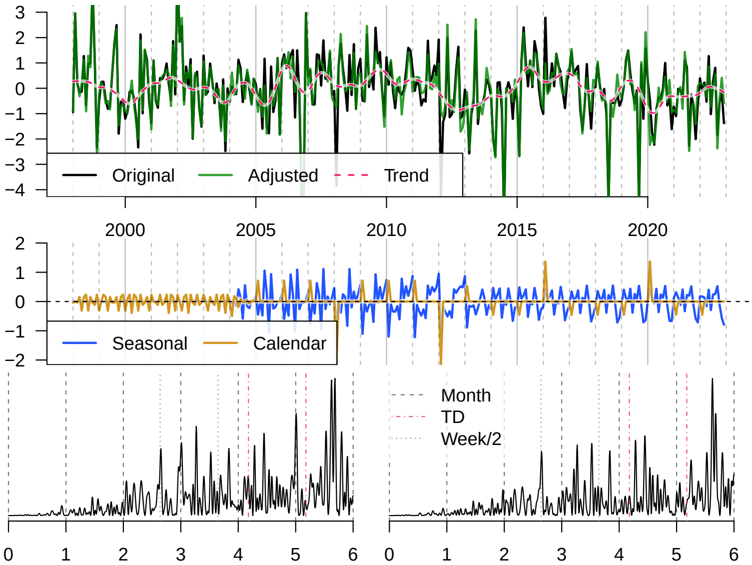

Visual analysis of multiple series

- The example above was **pure white noise**!
  - $\hat{S}_t$ looks very noisy
  - Period-wise plot shows no fluctuations
  - No seasonal spectral peaks
  - Full-sample $M_7 = 1.5$ (robust) or 2.3 (non-robust)
- Better pre-test and visualise long series, and choose a **single strategy**
  - E. g. fix transformation, TD, LY, Easter, allow only the RegARIMA model to change
- If the series X_t, Y_t, Z_t are similar in nature, consider pre-testing all and applying a single strategy

Seasonal adjustment in R

Pre-requisites

- R language (as of 2024-04, v. 4.3.3) and RStudio IDE
  - In case compilation is required in the future, RTools
- Java Development Toolkit (as of 2024-04, v. 22.0.1)
- JDemetra+ (as of 2024-04, v. 3.2.2)
- Within R, the packages `RJDemetra` (official), `rjd3toolkit` and `rjd3x13` (explained in the next slides)
  - The latter two are under active development, but they solve the problems that the GUI version of JDemetra+ has

Seasonal adjustment in R

- Use `RJDemetra` (official CRAN repository), `rjd3x13` and `rjd3toolkit` (GitHub, curated by Eurostat)
- Do **not** perceive the diagnostic statistics as the truth
  - The eye + economic reasoning are more important
  - Use wrappers for plots: original vs. adjusted, seasonal and calendar effects, spectra before-after, sub-series
  - Use HAC variance estimators in tests with a sane bandwidth
  - Run multiple versions of tests (e. g. without and with ranks), **be ready to defend your point under scrutiny**
- Rule #1: no seasonality $\Rightarrow$ no adjustment 'just in case'
- Export summaries (auto-selected specification, outlier dates and types, estimation period, robust test results)

RJDemetra and cousins

A fast and dynamically growing family of packages.

```
install.packages(c("RJDemetra", "remotes"))
library(remotes)
install_github("rjdemetra/rjd3toolkit@*release")
install_github("rjdemetra/rjd3x13@*release")
```

Pros: new convenience wrappers (custom calendars are a piece of cake now). **Cons:** the syntax may change; expected input for `RJDemetra` commands differs from that of `rjd3x13`; same goes for the output.

Feel free to experiment with `rjd3tramoseats`.

Extensive documentation for JD+ v3:
`jdemetra-new-documentation.netlify.app`

Write wrappers for custom arguments

- Create several national calendars with `rjd3toolkit::national_calendar()`
- Allow custom start and end of the estimation sample
- Automate cutting and blending of long series
  - Or allow custom estimation sub-samples and blending spans
- Auto-generate plots, force the user to examine them

Most useful wrappers

- Convert a data frame with a date variable to `ts` class (auto-detect the start and frequency)
- Robust seasonality and stable seasonality tests (to get a consistent M_7) based on SI_t
  - Take into account additive / multiplicative transformation
- Plot TS decomposition and add spectral diagrams pre- and post-adjustment
- Cut long series into chunks of equal sizes and blend multiple adjustments
  - Analyse how different the two adjustments are in the overlapping regions
- Reduce the full adjustment into a line of text to save for archival purposes (logging)

Consider allowing overriding arguments

The aforementioned parameters should be flexible:

- Transformation: log, none, or **auto**
- Trading days: 6, 1, none, or **auto**
  - Include initially for pre-treatment, then test the effect of removal – or do not include first and test addition
- Leap year and/or Easter: yes, no, or **auto**
  - Same: use or do not use in pre-treatment
- User-defined outliers

Use case 1: French calendar for SARIMAX

```r
s_d <- special_day  # A shortcut
cal.FR <- national_calendar(
  days = list(fixed_day(7, 14), # Bastille
    s_d('NEWYEAR'), s_d('CHRISTMAS'),
    s_d('MAYDAY'), s_d('EASTERMONDAY'),
    s_d('ASCENSION'), s_d('WHITMONDAY'),
    s_d('ASSUMPTION'), s_d('ALLSAINTSDAY'),
    s_d('ARMISTICE'))) # Reuse at any time
```

Given $\{Y_t\}$, generate calendar regressors for RegARIMA:

```r
WD.FR <- calendar_td(calendar = cal.FR, s = y,
           groups = c(1, 1, 1, 1, 1, 0, 0))
forecast::auto.arima(y, xreg = ycalWD, trace = TRUE)
forecast::auto.arima(y, xreg = NULL, trace = TRUE)
# If identical ARIMA pdq, compare the AICs
```

Use case 2: FR heating diesel HT price

Diagnose seasonality in heating diesel prices in France.
Enforce 6 TD or 1 WD, compare the final models.

```
y <- d[, "P_DIESCHAUF_HT_FR"]
cntxt <- modelling_context(
           calendars = list(FR = cal.FR))
spTD <- set_tradingdays(x13_spec("rsa5c"),
            option = "TradingDays", test = "None")
spWD <- set_tradingdays(x13_spec("rsa5c"),
            option = "WorkingDays", test = "None")
ysaTD <- x13(y, spTD, context = cntxt)
ysaWD <- x13(y, spWD, context = cntxt)
round(unlist(ysaWD$result$mstats), 2) # Bad?
plot(do.call(ts.union, sa_decomposition(ysaWD)))
```

Use case 2: FR heating diesel HT price

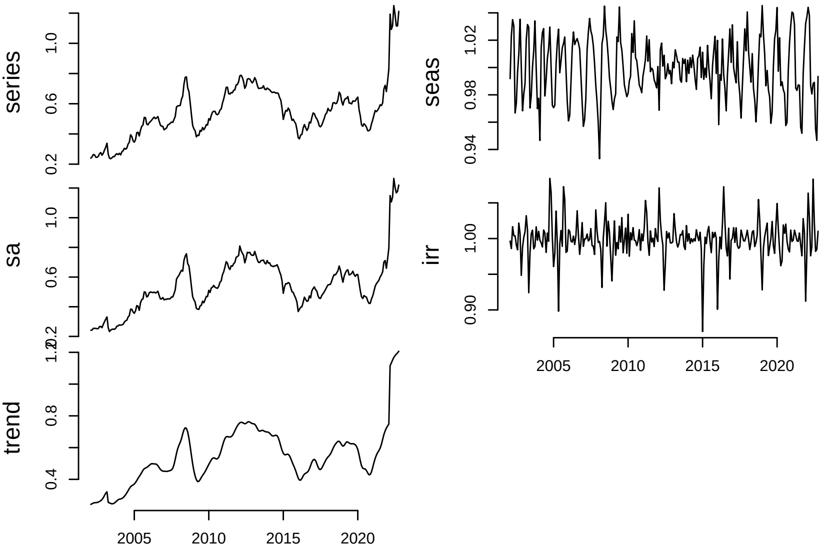

Use case 3: robust seasonality tests

- Presence of seasonality (via period dummies)
- Stability of seasonality (via yearly dummies)
  - Given in `example-04-seasonal-adjustment.R`
- Stability of the magnitude of $|\hat{S}_t|$
  - $\mathcal{H}_0$: the magnitude / variance of S_t is constant, $\mathcal{H}_1$: the magnitude of S_t depends on t (linear, spline etc.)
  - Given in `example-05-underestimating-seas.R`

A wrapper for the output of Census X13 binaries (for the users of `library(seasonal)`) is available in `helper-SA.R` – see `robustSeasTests()`.

Can you write a similar wrapper for `rjd3toolkit`'s `sa_decomposition(x13(...))`? **(Homework option!)**

Thank you for your attention!