

Contents

1	Introduction	1
2	Data	2
2.1	Absolute Data	2
2.2	Transformed Data	5
3	Methodology	9
3.1	Seasonal Autoregressive Integrated Moving Average (SARIMA)	9
3.2	Exponential Smoothing State Space Model (ETS)	10
4	Results	11
4.1	Seasonal Autoregressive Integrated Moving Average (SARIMA)	11
4.2	Exponential Smoothing State Space Model (ETS)	12
4.3	Time Series Cross-Validation Accuracy	13
5	Conclusions	15
	List of references	I
6	Appendix	III

1 Introduction

Decision-making institutions are required to deal with a high degree of uncertainty about the future development of an economy, but also about its current state (Walker & Marchau, 2003). Gross domestic product (GDP) is an indicator of economic activity that can be used to measure and compare levels of economic development of individual countries (Indergand & Leist, 2014). The uncertainty about the current level of economic activity derives from the time-lagged availability of quarterly GDP data (Siliverstovs, 2011). In Switzerland, the annual national accounts (ANA) are provided by the Federal Statistical Office (FSO), whereas quarterly national accounts (QNA) are issued by the State Secretariat for Economic Affairs (SECO) (Indergand & Leist, 2014). The Swiss QNAs are generated by an indirect estimation methodology, which relies on temporal disaggregation techniques (Siliverstovs, 2011). Quarterly estimations are

deduced from available annual data, using both statistical procedures and quarterly indicators (Indergand & Leist, 2014). The accuracy of the preliminary annual estimates, which are the sum of the quarterly estimates, depends heavily on the quality of the available and retained quarterly indicators (Siliverstovs, 2011). The time lag of GDP data publication in Switzerland adds up to approximately two months, which is particularly high compared to large economies such as the U.S. (about one month) and European countries (about six weeks) (Ruenstler et al, 2009; Siliverstovs, 2011). In addition, initial releases of GDP data are frequently being revised by statistical agencies whenever more comprehensive information becomes available (Boysen-Hogrefe & Neuwirth, 2012).

An accurate forecast of the GDP is a useful way to get an indication of the general direction of economic activity in the future (Islam & Clarke, 2001). Its relevance for national governments emerges as GDP development has implications for economic development strategies, economic policies and the allocation of funding (Baffigi & Golinelli, 2004). Therefore, this paper provides an approach to offer a solid GDP forecast for the use case of Switzerland. At the beginning of this paper, the selected time series will be presented and described. After having presented the data, the model selection and methodology will be elaborated in the following chapter. Next, the selected forecast methods will be performed in order to create real GDP forecasts for Switzerland. Subsequently, the results will be discussed and at the end of this paper, a conclusion that provides a summary of the findings will be given.

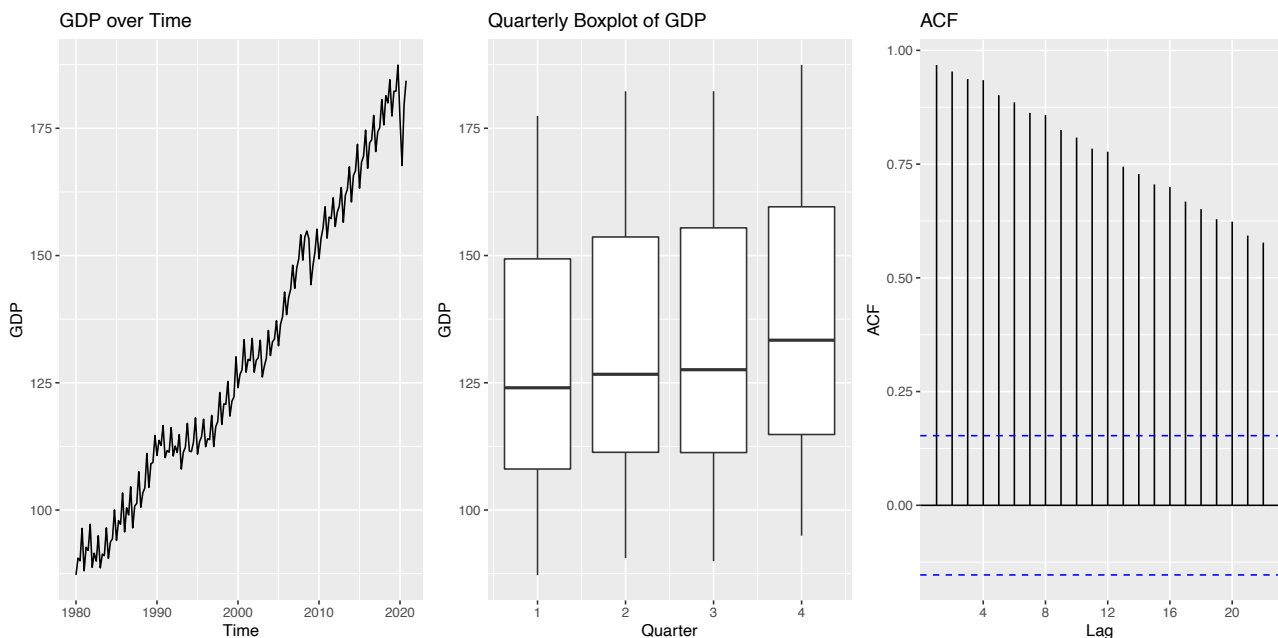
2 Data

2.1 Absolute Data

In this subsection, the absolute level of GDP is analysed. To provide an overview of the data, descriptive insights will be presented. Then, statistical unit root tests such as the Augmented-Dickey-Fuller (ADF) and the Kwiatkowski, Phillips, Schmidt and Shin (KPSS) are implemented in order to check stationarity and to detect possible trend patterns of the underlying data.


The quarterly GDP variable has constantly increased over the analysed time period between 1980 and 2020. The 164 data observations vary between 87.3 and 187.4, whereas the median amounts to 126.9 and the mean to 130.9. The average quarterly GDP growth rate of Switzerland between 1980 and 2020 (compared to previous year value) amounts to 1.7 percent. In addition, comparatively clear seasonality patterns can already be detected. While the averages of quarter 2 and 3 are relatively close to each other, it can be stated that the mean of quarter 1 is clearly below, whereas the mean of quarter 4 is clearly above the others (compare boxplot below). This also indicates seasonality in the data even when considering the general growth trend.

The dotted lines of the correlogram (ACF plot below) illustrate a 95 percent interval. Therefore, for a large number (for all illustrated lags) autocorrelation of lags is highly significant as the ACF values are above the dotted line. This is also reflected by the fact that the Box-Pierce test (appendix 1) rejects the null of no-autocorrelation. So, the evidence supports that the absolute level of GDP is autocorrelated. Also, the ACF value for the lag of period 1 is really close to one. Nonetheless, the values of ACF are clearly decreasing over time (trend), however, still remaining on highly significant levels.



Looking at the illustrations of the decomposition (see decomposition illustration below), the trend and seasonality information extracted from the time series seem to be reasonable.

Multiplicative decompositions are considered as standard when dealing with economic time series on an absolute level (Hyndman & Athanasopoulos, 2021). The seasonal pattern in the seasonal plot seems to be very simple. Simpler patterns are usually more intuitive to model and lead to more accurate predictions (Hyndman & Athanasopoulos, 2021). The trend seems to be more or less strictly positive over time and could be exponential. The residuals are also interesting, showing almost similar variability for all periods over time. However, a potential outlier in 2020 (Q1) can be detected, which seems appropriate due to the coronavirus effect. Overall, the variation is considered as extreme (as e.g., in strong heteroscedasticity).



Due to the possible presence of a trend, the right model for the KPSS and ADF test is a random walk with drift and trend (Hyndman & Athanasopoulos, 2021). Therefore, the type “tau” in KPSS and “trend” in ADF have been selected. Considering both a unit root test (ADF) as well as a stationarity test (KPSS) allows reducing the chance of falsely rejecting the null hypothesis (Kwiatkowski et al, 1992). For the underlying data, the t-value of the KPSS test (tau with 4 lags, appendix 2) is higher than the significance level. Due to these circumstances, the null of stationarity can be rejected and therefore, the KPSS suggests non-stationary data. In the ADF test (appendix 3-4), the t-stat for tau3 is lower than its corresponding significance level, therefore the null cannot be rejected and the ADF test suggests that the variable has as a unit

root. The second t-statistic (ϕ_2) is larger than the significance level at 5 percent, however, not larger than the significance level at 1 percent. Still, the null, suggesting the presence of unit root with drift and trend, can be rejected. ϕ_3 is smaller than the significance level, thus the null (presence of unit root and trend) cannot be rejected. To conclude, the time series can be described as non-stationary data with drift. Further ADF and KPSS test results respect this observation (again on 5 percent significance level, appendix 5-6).

Also, as the distribution of the residuals shows (see appendix 7), the residuals do not really follow a normal distribution. Therefore, a data transformation makes sense. E.g., a log transformation would be appropriate, as it is often used for exponential phenomena like economic time series (Luetkepohl & Xu, 2009). In addition, the log transformation is widely used, makes parameter better interpretable and can save one differentiation (Fink, 2009). However, a data transformation is also always closely related to a loss of information (Luetkepohl & Xu, 2009). Nonetheless, differencing might help to stabilize the mean of the time series and might be able to reduce trend and seasonal patterns (Hyndman & Athanasopoulos, 2021).


2.2 Transformed Data

The subsection above has demonstrated that further data transformation procedures are required. Therefore, this subsection will implement transformation procedures such as log-transformation and differencing. Subsequently, ADF and KPSS tests are elaborated again in order to check if the transformation process has converted data to stationary and also to analyse trend patterns. Then, a brief intuitive analysis of the ACF and PACF will follow. At the end of this subsection, a test for structural break will be performed.

After log-transforming and differencing, an important point consists of the fact that the log differences of two time series points directly return growth rates (Hyndman & Athanasopoulos, 2021). Log-differencing is really helpful and makes investigations more interpretable and results in clearer solutions (Fink, 2009). In addition, it also provides a sort of normalization, so, the data looks less skewed after having used the transformation (Hyndman & Athanasopoulos,


2021). In a first step, the log-transformation is performed and the unit root tests are implemented (appendix 8-11). When analysing the log-transformed data via the ADF test with parameter trend, a clear deterministic trend cannot be found (appendix 8-9). However, data has not become stationary and is still dealing with a significant drift (appendix 10-11). After changing the test parameter from trend to drift, results remained the same as before. Also, KPSS test results respect this observation.

Looking at the illustrations of the decomposition of the log-transformed data (see decomposition illustration below), the trend and seasonality information extracted from the time series have become a bit more clear. Additive decompositions are considered as standard for handling log-transformed data (Hyndman & Athanasopoulos, 2021)



As already mentioned in the previous subsection, differencing might help to stabilize the mean of the time series and could be able to reduce trend and seasonal patterns (Hyndman & Athanasopoulos, 2021). In addition, seasonal differencing will remove seasonal effects (Holmes, Scheuerell & Ward, 2021). As the underlying data refers to quarters, the seasonal-differencing will be performed in a first-order with lag 4 (Hyndman & Athanasopoulos, 2021). Also, as a deterministic trend was not found for the log-transformed data before (appendix 8-11) and the series does not show any signs of a clear trend (see decomposition illustration below), the tests


for the log-transformed data with differencing were directly elaborated from the starting point of a random walk with drift. So, in this use case, after log-transforming and differencing, the data became stationary as the t-value of τ_3 in the ADF test is clearly above the significance level (ADF test results, appendix 12-13). Therefore, the null is rejected and a unit root cannot be found. However, the data is still dealing with a drift as ϕ_2 is above the respective significance level. Also, the results of the KPSS test clearly confirm the above determination as the null cannot be rejected (KPSS test result, appendix 14).



Next, a descriptive analysis of ACF and PACF is following. PACF finds correlation of the residuals, after removing the effects which are already explained by the earlier lags, for the next lag value (partial instead of complete) (Hyndman & Athanasopoulos, 2021). So, if there exists any hidden information in the residuals which can be modeled by the next lag, it results in a high correlation. For the underlying data, the ACF changes sign between the first and the sixth lag and then it remains negative afterwards. The PACF presents one highly significant spike in period one and one scarcely significant spike in period three. Further explanations of the discussed behavior of this process will be provided in the following chapter of methodology in which also the model selection will be discussed.

When looking at the residuals (appendix 15), it can be stated that the standard deviation is

not perfectly constant with time. The residuals do not perfectly look like white noise because. There still seems to be information remaining in the residuals which has not been exactly captured. As the distribution of the residuals shows, the distribution looks in tendency like a normal distribution. However, it does not fit perfectly well.



When elaborating prediction models, structural breaks can usually cause forecasting errors and lead to unreliable models. Therefore, the ordinary-least-squares (OLS) based F-statistics test can be used in order to find structural breaks in time series (Hanck et al, 2020). The F-statistics test further extends the Chow test via calculating the F statistics for all potential change points in an interval (Zeileis et al, 2002). Subsequently, the test rejects the null if one of these statistical values get too large (Zeileis et al, 2002, p. 10-11). The black line in the plot (appendix 16) is the set of F-statistics. The maximum Fstats is the quant likelihood ratio (QLR) statistics (Hanck et al, 2020). The red line represents the critical values. In this case (see appendix 16), the null cannot be rejected and therefore there the test suggests no structural change. When looking at the p-value of the sc test (null-hypothesis of no structural breaks), the null cannot be rejected, therefore the data seems not to be dealing with structural breaks.

3 Methodology

Exponential smoothing and ARIMA models are two of the most popular approaches when it comes to time series forecasting (Namin & Namin, 2018). While ARIMA models try to describe autocorrelation procedures, exponential smoothing models focus on the trend and seasonality description in the data (Hyndman & Athanasopoulos, 2021). Therefore, exponential smoothing and ARIMA models, as widely-used time series forecasting methods, provide complementary approaches. In this section, the seasonal autoregressive integrated moving average (SARIMA) model and exponential smoothing state space (ETS) model will be briefly explained and subsequently, the model selection will be discussed.

3.1 Seasonal Autoregressive Integrated Moving Average (SARIMA)

In econometrics, seasonal time series forecasting is traditionally performed via SARIMA models, which are considered as one of the most popular methods (Choi, Yu & Au, 2011). SARIMA models are able to use seasonal differencing in order to remove non-stationarity from time series (Adhikari & Agrawal, 2013). The AR part in the model connotes “autoregressive”, the I stands for “integrated”, and the MA for “moving-average” (Holmes, Scheuerell & Ward, 2021). The term “autoregressive” implicates a regression of the variable of interest against itself (Hyndman & Athanasopoulos, 2021). So, autoregressive models use linear combinations of past values of the variable of interest. The “integrated” part indicates the number of times that data observations have been replaced with the difference between their values and the previous values (Hyndman & Athanasopoulos, 2021). The “moving-average” (MA) process is represented by the weighted sum of the current random error plus the most recent errors (Holmes, Scheuerell & Ward, 2021).

Based on the ACF and PACF plots, it is proposed to only include the lags with the most significant lags in order to enable a parsimonious model (Adhikari & Agrawal, 2013). Due to the fact that the PACF plot only includes two relatively clear significant lag, a possible AR(2) term is suggested. As the ACF-plot shows sinusoidal behavior (Hyndman & Athanasopoulos,

2021) and the other values are barely significant, an MA(0) is selected. Therefore, the suggested order for the non-seasonal part is (2,0,0). Although the section above has demonstrated that seasonal patterns seem to exist, a clear suggestion for the non-seasonal part of the SARIMA model cannot be projected precisely. Hence, the auto.arima is used. The function returns the best fitting ARIMA model according to either AIC (Akaike information criteria), AICC (AIC adjusted for small samples) or BIC (Bayesian information criteria) information criteria value (Hyndman & Athanasopoulos, 2021). It conducts a search over all possible models within the order constraints that were provided (Hyndman & Athanasopoulos, 2021). For the GDP growth rates of Switzerland, auto.arima suggests the use of an ARIMA with the order (2,0,1)(0,0,1)[4] (see summary of auto.arima for log-diff data in appendix 17). When only using log-transformed data as input, auto.arima returns the same order (2,0,1)(0,1,1)[4], however, also including a seasonal difference (see summary of auto.arima for log-transformed data in appendix 18).

3.2 Exponential Smoothing State Space Model (ETS)

ETS forecasting is an approach for predicting time series (Jofipasi et al, 2018). The flexibility of the ETS model lies in its ability to make use of trend and seasonal components of different characteristics (Hyndman & Athanasopoulos, 2021). So, via a simple integrated approach, exponentially weighted moving averages allow estimating both trend and seasonal effects (Holt, 2004). These ETS models are called state space models (Hyndman & Athanasopoulos, 2021). Instead of minimising the sum of squared estimate of errors, the parameters in ETS are estimated via maximizing the likelihood (Hyndman & Athanasopoulos, 2021). Good models are usually characterized by large likelihoods, as this represents a large probability of the data arising from the specified model (Hyndman & Athanasopoulos, 2021). So, the initial states (l, b, s), so-called starting values, and the smoothing parameters (alpha, beta and gamma) are estimated in order to maximize the likelihood (Hyndman & Athanasopoulos, 2021). The parameter alpha specifies the coefficient of the level smoothing, beta specifies the coefficient of trend smoothing and gamma specifies the coefficient of the seasonal smoothing (Hyndman & Athanasopoulos, 2021). The initial state l represents the smoothed estimate of level at time

t , whereas b is the smoothed estimate of change in the trend in time t and s stands for the smoothed estimate of the appropriate seasonal component at time t (Emmanuel et al, 2014).

Given the time series of the GDP of Switzerland, the methodological approach of ETS suits very well as the previous discussion from above has demonstrated that the data is confronted with a drift and seasonality. Due to the fact that the time series has been log-transformed, an ETS with additive error type, additive trend type and additive season type seems recommendable. So, the suggested model order is (A,A,A). Compared to the model suggested, the automatically selected version (Z,Z,Z) suggests using a model with a multiplicative error term, an additive trend type and a multiplicative seasonal term (M,A,M). As the model performance improved (appendix 19-20), the final selection is the so-called multiplicative Holt-Winters method with multiplicative errors (M,A,M).


4 Results

After having discussed the main characteristics of the selected SARIMA model and ETS model, this section will discuss the individual performance and the resulting parameters of the before-selected SARIMA and ETS model for the log-transformed data. At the end of this section, the accuracy of both elaborated and more sophisticated models will also be compared to each other and to simple approaches (for benchmarking reasons) such as the naive method, the mean approach and the drift method.

4.1 Seasonal Autoregressive Integrated Moving Average (SARIMA)

The selected ARIMA(2,0,1)(0,1,1)[4] with drift respects the above-made observation that the data is dealing with a drift. Compared to the level of the coefficients, the training MAE and RMSE of the coefficients seem to be quite low. The residuals for the fitted model are shown in the figure below. The residuals of the model perform relatively well as they are uncorrelated and their distribution does not seem extremely skewed. However, the model presents a few outliers (e.g., financial crisis 2008 and coronavirus crisis 2020) which might

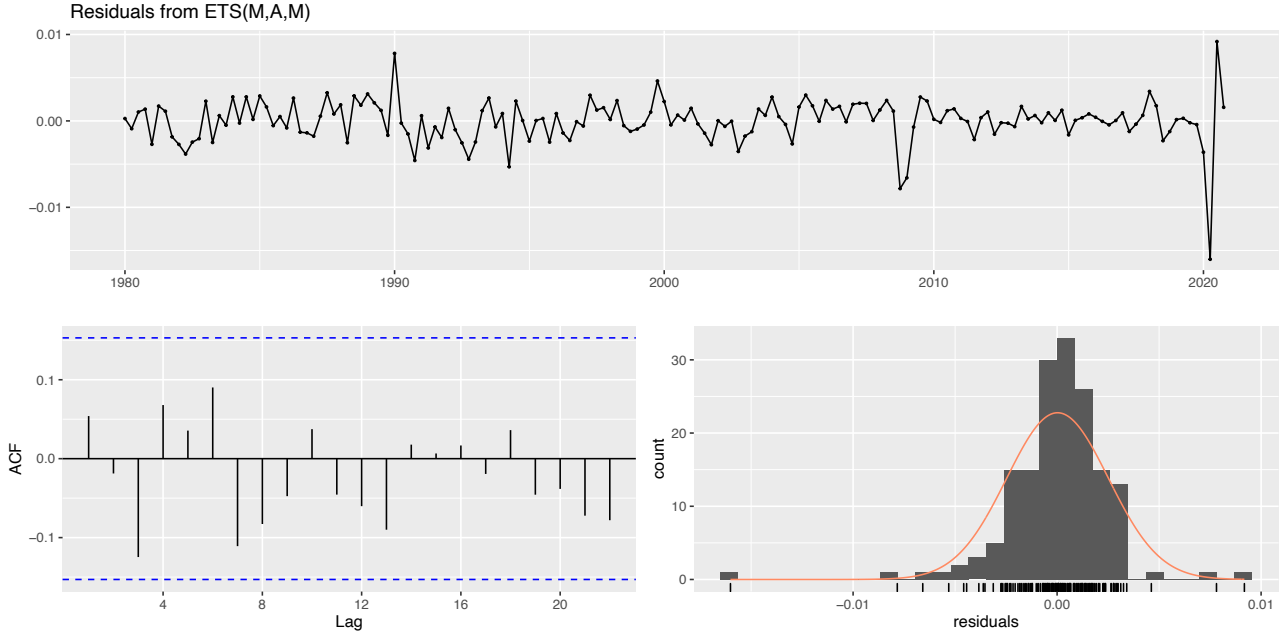
indicates heteroskedasticity and fat tails. Nonetheless, all in all, it can be stated that it seems like the model is able to incorporate a large share of the dynamics of the residuals. The Ljung-Box test returns a relatively large p-value, also suggesting that the residuals are white noise (Hyndman & Athanasopoulos, 2021).



4.2 Exponential Smoothing State Space Model (ETS)

For the time series of the GDP of Switzerland, the so-called multiplicative Holt-Winters method with multiplicative errors (M,A,M) persuades via relatively low training error measures. The smoothing parameter alpha amounts to approximately 0,75 which indicates that more weight to relatively new data observations will be added (Hyndman & Athanasopoulos, 2021). The value of beta is really close to zero and therefore, the estimate of the slope of the trend component is almost not updated over time (Emmanuel et al, 2014). So, while the level changes over time, the slope of the trend component remains more or less the same and is almost set equal to its initial value. The rather low gamma value (0,22) indicates that the estimate of the seasonal component is based on both recent and less recent observations (Emmanuel et al, 2014). Next, the residuals of the model perform relatively well as they are uncorrelated and their distribution does not seem extremely skewed. However, the model presents a few extreme outliers which

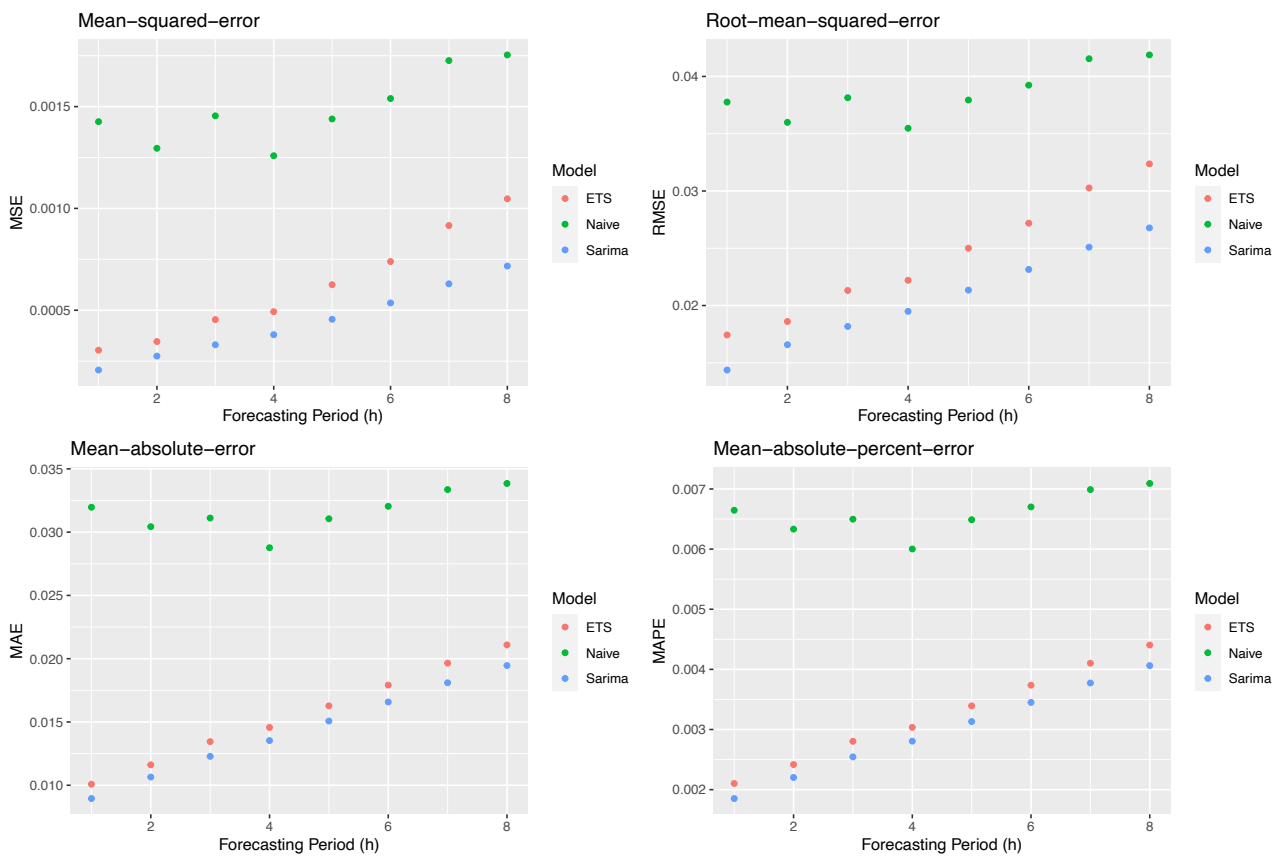
might indicate heteroskedasticity and fat tails. The largest outliers resulted from the financial crisis (2008) and from the coronavirus effect (2020) and can therefore be justified. Nonetheless, all in all, it can be stated that it seems like the model is able to incorporate a large share of residuals dynamics.



4.3 Time Series Cross-Validation Accuracy

In order to check the accuracy of the above-proposed models and the simple approaches, a rather more sophisticated method, so-called cross-validation, was performed. Cross-validation, as statistical technique, splits the whole dataset into small train and test chunks and therefore it allows to include the entire set of data (Burger, 2018). After having evaluated the error for each chunk, those final errors are averaged (Burger, 2018). When it comes to time series analysis, the process is also called evaluation on a rolling forecasting origin as the origin at which the forecast is based rolls forward with regards to time (Hyndman & Athanasopoulos, 2021). The illustration below evaluates the forecasting performance of a one to eight-step-ahead forecast. When comparing MAE, MSE, RMSE and MAPE of the models with regards to the unseen test data (via cross-validation), the more sophisticated models (SARIMA and ETCS) performed significantly better than the simple approaches. Due to comparability, the figure excludes

two simple approaches (mean and drift), as they perform significantly worse than the other methodological approaches. However, all approaches are included in the illustration in appendix 21. Within the simple approaches, the naive methods worked best, followed by the mean method and the drift method (appendix 21). Still, according to forecast accuracy measures, the SARIMA model worked best as forecast method for the underlying data (comparing selected relevant performance measures in the plot below). However, the performance measure difference of SARIMA, compared to ETS, is not considered as large. Previous studies also confirm this observation as they indicate that Winters exponential smoothing and ARIMA both perform well when the macroeconomic environment is relatively stable over the time horizon (Ramos, Santos & Rebelo, 2015). In addition, the plot below also demonstrates that the forecast error mostly increases with an increasing forecast horizon, which makes perfectly sense.



5 Conclusions

In this paper, a seasonal autoregressive integrated moving average (SARIMA) model and an exponential smoothing state space (ETS) model for the use case of the GDP of Switzerland were elaborated. This paper contributes to the relevance of previous practical research using historical data to forecast GDP in Switzerland. When evaluating the performance measures of the final models via cross-validation, it can be concluded that both perform significantly better than just simple approaches such as the naive, mean or drift method. According to forecast performance measures, the SARIMA model worked best as forecast method for the underlying data, however, the performance measure differences of SARIMA, compared to ETS, is not extremely dissimilar. Also previous studies (Ramos, Santos & Rebelo, 2015) confirm these results. With regards to the simple approach methods, the drift and mean approaches did not perform well, while the naive method worked quite fair. Still, when comparing the performance of more sophisticated models in the short and in the longer run via cross-validation, these models performed significantly better.

List of references

- Adhikari, R. & Agrawal, R. (2013). An Introductory Study on Time series Modeling and Forecasting.
- Baffigi, A. & Golinelli, R. (2004). Bridge models to forecast the Euro area GDP. *International Journal of Forecasting*, 20(3), 447-460.
- Boysen-Hogrefe, J. & Neuwirth, S. (2012). The impact of seasonal and price adjustments on the predictability of German GDP revisions. *Kiel Working Paper*, 1753.
- Burger, S. (2018). Introduction to Machine Learning with R. Sebastopol: O'Reilly.
- Choi, T., Yu, Y. & Au, K. (2011). A hybrid SARIMA wavelet transform method for sales forecasting. *Decision Support Systems*, 51(1), 130-140.
- Fink, E. (2009). The FAQs on Data Transformation. *Communication Monographs*, 76(4), 379-397.
- Hanck, C., Arnold, M., Gerber, A. & Schmelzer, M. (2020). Introduction to Econometrics with R. Retrieved June 3, 2021 from <https://www.econometrics-with-r.org>
- Holmes, E, Scheuerell, M. & Ward, E. (2021). Applied Time Series Analysis for Fisheries and Environmental Sciences. Retrieved June 3, 2021 from <https://nwfs-timeseries.github.io/atsa-labs/>
- Holt, C. (2004). Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting*, 20(1), 5-10.
- Hyndman, R. & Athanasopoulos, G. (2021). Forecasting: Principles and Practice. Retrieved June 3, 2021 from <https://otexts.com/fpp2/>
- Indergand, R & Leist, S. (2014). A Real-Time Data Set for Switzerland. *Swiss Journal of Economics and Statistics*, 150(4), 331-352.
- Islam, S. & Clarke, M. (2002). The Relationship between Economic Development and Social

- Welfare: A New Adjusted GDP Measure of Welfare. *Social Indicators Research*, 57(2), 201-228.
- Jofipasi, C., Miftahuddin, M. & Sofyan, H. (2018). Selection for the best ETS (error, trend, seasonal) model to forecast weather in the Aceh Besar District. *IOP Conference Series: Materials Science and Engineering*, 352(1).
- Kwiatkowski, D., Phillips, P., Schmidt, P. & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, 54, 159-178.
- Luetkepohl, H. & Xu, F. (2009). The Role of the Log Transformation in Forecasting Economic Variables. *CESifo Working Paper*, 2591.
- Ramos, P., Santos, N. & Rebelo, R. (2015). Performance of state space and ARIMA models for consumer retail sales forecasting. *Robotics and Computer-Integrated Manufacturing*, 34, 151-163.
- Ruenstler, G., Barhoumi, K., Benk, S., Cristadoro, R., den Reijer, A., Jakaitiene, A., Jelonek, P., Rua, A., Ruth, K., & Nieuwenhuyze, C. (2009). Short-Term Forecasting of GDP Using Large Datasets: A Pseudo Real-Time Forecast Evaluation Exercise. *Journal of Forecasting*, 28, 595-611.
- Siarni Namini, S. & Siarni Namin, A. (2018). Forecasting Economics and Financial Time Series: ARIMA vs. LSTM. Retrieved June 3, 2021 from https://www.researchgate.net/publication/323867492_Forecasting_Economics_and_Financial_Time_Series_ARIMA_vs_LSTM
- Siliverstovs, B. (2011). Are GDP revisions predictable? Evidence for Switzerland. *KOF Working Papers*, 281.
- Walker, W. & Marchau, V. (2003). Dealing With Uncertainty in Policy Analysis and Policy-making. *Integrated Assessment*, 4(1), 1-4.
- Zeileis, A. Leisch, F., Hornik, K. & Kleiber, C. (2002). Strucchange: An R package for testing for structural change in linear regression models. *Journal of Statistical Software*, 7(2).

6 Appendix

1) Box-Pierce-test for absolute data

```
##  
## Box-Pierce test  
##  
## data: myts  
## X-squared = 153.59, df = 1, p-value < 2.2e-16
```

2) KPSS-test for absolute data (type: tau)

```
##  
## #####  
## # KPSS Unit Root Test #  
## #####  
##  
## Test is of type: tau with 4 lags.  
##  
## Value of test-statistic is: 0.6108  
##  
## Critical value for a significance level of:  
##          10pct  5pct 2.5pct  1pct  
## critical values 0.119 0.146 0.176 0.216
```

3) ADF-test for absolute data (type: trend, lags: 10)

```
##  
## #####  
## # Augmented Dickey-Fuller Test Unit Root Test #  
## #####  
##
```

```

## Test regression trend

##

##

## Call:
## lm(formula = z.diff ~ z.lag.1 + 1 + tt + z.diff.lag)
##

## Residuals:
##      Min       1Q   Median       3Q      Max
## -13.9363  -0.6513   0.1193   0.8383   8.3280
##

## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)   6.958463   3.364264   2.068 0.040448 *
## z.lag.1       -0.077135   0.042163  -1.829 0.069463 .
## tt            0.047529   0.024917   1.907 0.058509 .
## z.diff.lag1   -0.195648   0.087535  -2.235 0.026996 *
## z.diff.lag2   -0.170488   0.097374  -1.751 0.082160 .
## z.diff.lag3   -0.278234   0.130140  -2.138 0.034257 *
## z.diff.lag4    0.455819   0.130554   3.491 0.000643 ***
## z.diff.lag5   -0.009262   0.135613  -0.068 0.945647
## z.diff.lag6   -0.016092   0.137381  -0.117 0.906920
## z.diff.lag7   -0.090683   0.130638  -0.694 0.488737
## z.diff.lag8    0.139429   0.128445   1.086 0.279559
## z.diff.lag9   -0.100123   0.127057  -0.788 0.432020
## z.diff.lag10  -0.177433   0.129964  -1.365 0.174365
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##

```

```
## Residual standard error: 1.914 on 140 degrees of freedom
## Multiple R-squared:  0.8451, Adjusted R-squared:  0.8318
## F-statistic: 63.64 on 12 and 140 DF,  p-value: < 2.2e-16
##
##
## Value of test-statistic is: -1.8294 5.616 1.8865
##
## Critical values for test statistics:
##      1pct  5pct 10pct
## tau3 -3.99 -3.43 -3.13
## phi2  6.22  4.75  4.07
## phi3  8.43  6.49  5.47
```

4) ADF-test for absolute data (type: trend, lags: 4)

```
##
## #####
## # Augmented Dickey-Fuller Test Unit Root Test #
## #####
##
## Test regression trend
##
##
## Call:
## lm(formula = z.diff ~ z.lag.1 + 1 + tt + z.diff.lag)
##
## Residuals:
```

	Min	1Q	Median	3Q	Max
##	-13.7893	-0.9357	0.2397	0.8222	7.6883

```

##
## Coefficients:
##           Estimate Std. Error t value Pr(>|t|)
## (Intercept)  8.39868    2.97202   2.826 0.005349 **
## z.lag.1      -0.09643    0.03631  -2.655 0.008763 **
## tt           0.05727    0.02159   2.653 0.008826 **
## z.diff.lag1  -0.18439    0.06787  -2.717 0.007356 **
## z.diff.lag2  -0.22756    0.06353  -3.582 0.000459 ***
## z.diff.lag3  -0.25621    0.06608  -3.877 0.000157 ***
## z.diff.lag4   0.68490    0.06738  10.164 < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 1.889 on 152 degrees of freedom
## Multiple R-squared:  0.8424, Adjusted R-squared:  0.8362
## F-statistic: 135.4 on 6 and 152 DF,  p-value: < 2.2e-16
##
##
## Value of test-statistic is: -2.6555 5.575 3.5434
##
## Critical values for test statistics:
##           1pct  5pct 10pct
## tau3 -3.99 -3.43 -3.13
## phi2  6.22  4.75  4.07
## phi3  8.43  6.49  5.47

```

5) ADF-test for absolute data (type: drift, lags: 4)

```

##

```

```

## #####
## # Augmented Dickey-Fuller Test Unit Root Test #
## #####
##
## Test regression drift
##
##
## Call:
## lm(formula = z.diff ~ z.lag.1 + 1 + z.diff.lag)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -14.2215  -0.8600   0.2148   0.9421   7.6629
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  0.758266   0.748459   1.013 0.312610
## z.lag.1      -0.001220   0.005651  -0.216 0.829383
## z.diff.lag1  -0.252660   0.064030  -3.946 0.000121 ***
## z.diff.lag2  -0.274936   0.062165  -4.423 1.84e-05 ***
## z.diff.lag3  -0.284666   0.066478  -4.282 3.26e-05 ***
## z.diff.lag4   0.680673   0.068679   9.911 < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 1.925 on 153 degrees of freedom
## Multiple R-squared:  0.8351, Adjusted R-squared:  0.8297
## F-statistic: 155 on 5 and 153 DF, p-value: < 2.2e-16

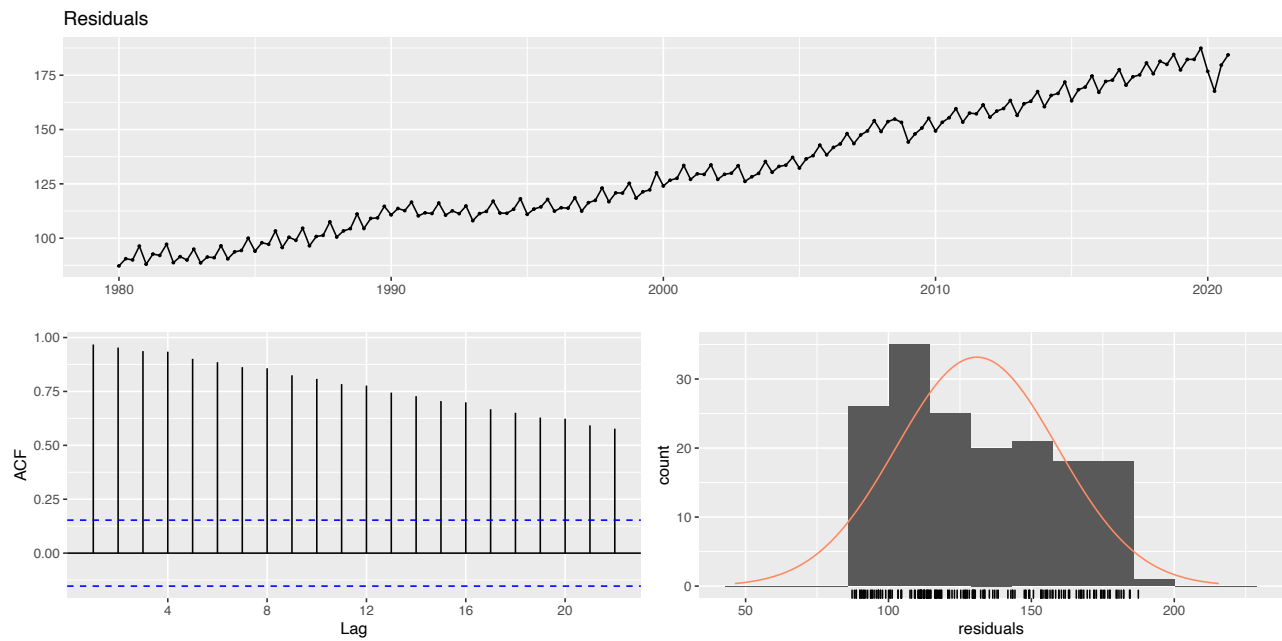
```

```
##
##
## Value of test-statistic is: -0.2159 4.6594
##
## Critical values for test statistics:
##      1pct  5pct 10pct
## tau2 -3.46 -2.88 -2.57
## phi1  6.52  4.63  3.81
```

6) KPSS-test for absolute data (type: mu)

```
##
## #####
## # KPSS Unit Root Test #
## #####
##
## Test is of type: mu with 4 lags.
##
## Value of test-statistic is: 3.3071
##
## Critical value for a significance level of:
##      10pct  5pct 2.5pct  1pct
## critical values 0.347 0.463  0.574 0.739
```

7) Checkresiduals for absolute data



8) ADF-test for log-transformed data (type: trend, lags: 10)

```
##
## #####
## # Augmented Dickey-Fuller Test Unit Root Test #
## #####
##
## Test regression trend
##
##
## Call:
## lm(formula = z.diff ~ z.lag.1 + 1 + tt + z.diff.lag)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -0.080908 -0.005201  0.000974  0.006158  0.050998
##
## Coefficients:
```



```

##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)   0.6507862  0.2559453   2.543   0.0121 *
## z.lag.1       -0.1442000  0.0574343  -2.511   0.0132 *
## tt            0.0006359  0.0002600   2.445   0.0157 *
## z.diff.lag1   -0.0896016  0.0893910  -1.002   0.3179
## z.diff.lag2   -0.0776105  0.0949173  -0.818   0.4149
## z.diff.lag3   -0.1717814  0.1098947  -1.563   0.1203
## z.diff.lag4    0.4821488  0.1087112   4.435 1.85e-05 ***
## z.diff.lag5    0.0131903  0.1153634   0.114   0.9091
## z.diff.lag6    0.0307891  0.1149592   0.268   0.7892
## z.diff.lag7   -0.0435746  0.1097339  -0.397   0.6919
## z.diff.lag8    0.2210422  0.1062769   2.080   0.0394 *
## z.diff.lag9   -0.0580834  0.1063606  -0.546   0.5859
## z.diff.lag10  -0.1392596  0.1075106  -1.295   0.1973
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.01259 on 140 degrees of freedom
## Multiple R-squared:  0.8889, Adjusted R-squared:  0.8794
## F-statistic: 93.32 on 12 and 140 DF,  p-value: < 2.2e-16
##
##
## Value of test-statistic is: -2.5107 7.2737 3.3415
##
## Critical values for test statistics:
##      1pct  5pct 10pct
## tau3 -3.99 -3.43 -3.13
## phi2  6.22  4.75  4.07

```

```
## phi3 8.43 6.49 5.47
```

9) KPSS-test for log-transformed data (type: tau)

```
##
```

```
## #####
```

```
## # KPSS Unit Root Test #
```

```
## #####
```

```
##
```

```
## Test is of type: tau with 4 lags.
```

```
##
```

```
## Value of test-statistic is: 0.1837
```

```
##
```

```
## Critical value for a significance level of:
```

```
##          10pct  5pct 2.5pct  1pct
```

```
## critical values 0.119 0.146 0.176 0.216
```

10) ADF-test for log-transformed data (type: drift, lags: 4)

```
##
```

```
## #####
```

```
## # Augmented Dickey-Fuller Test Unit Root Test #
```

```
## #####
```

```
##
```

```
## Test regression drift
```

```
##
```

```
##
```

```
## Call:
```

```
## lm(formula = z.diff ~ z.lag.1 + 1 + z.diff.lag)
```

```
##
```

```
## Residuals:
```

```

##           Min           1Q       Median           3Q           Max
## -0.081244 -0.006660  0.001860  0.007115  0.047767
##
## Coefficients:
##           Estimate Std. Error t value Pr(>|t|)
## (Intercept)  0.024323   0.024310   1.001 0.318640
## z.lag.1      -0.004168   0.005002  -0.833 0.405974
## z.diff.lag1 -0.223669   0.059007  -3.791 0.000216 ***
## z.diff.lag2 -0.226418   0.057777  -3.919 0.000134 ***
## z.diff.lag3 -0.244247   0.059778  -4.086 7.06e-05 ***
## z.diff.lag4  0.698888   0.060644  11.524 < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.01313 on 153 degrees of freedom
## Multiple R-squared:  0.8772, Adjusted R-squared:  0.8732
## F-statistic: 218.6 on 5 and 153 DF,  p-value: < 2.2e-16
##
##
## Value of test-statistic is: -0.8333 4.776
##
## Critical values for test statistics:
##           1pct   5pct 10pct
## tau2 -3.46 -2.88 -2.57
## phi1  6.52  4.63  3.81

```

11) KPSS-test for log-transformed data (type: mu)

```

##

```

```

## #####
## # KPSS Unit Root Test #
## #####
##
## Test is of type: mu with 4 lags.
##
## Value of test-statistic is: 3.3381
##
## Critical value for a significance level of:
##          10pct  5pct 2.5pct  1pct
## critical values 0.347 0.463 0.574 0.739

12) ADF-test for differences of log-transformed data (type: drift, lags: 10)

##
## #####
## # Augmented Dickey-Fuller Test Unit Root Test #
## #####
##
## Test regression drift
##
##
## Call:
## lm(formula = z.diff ~ z.lag.1 + 1 + z.diff.lag)
##
## Residuals:
##      Min      1Q  Median      3Q      Max
## -0.084200 -0.004258  0.001059  0.005690  0.050857
##

```

```

## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  0.008291  0.002264   3.663 0.000356 ***
## z.lag.1      -0.463739  0.111236  -4.169  5.4e-05 ***
## z.diff.lag1  0.261119  0.116404   2.243 0.026489 *
## z.diff.lag2  0.351441  0.126211   2.785 0.006119 **
## z.diff.lag3  0.188386  0.130218   1.447 0.150267
## z.diff.lag4 -0.197019  0.124068  -1.588 0.114594
## z.diff.lag5  0.212754  0.125897   1.690 0.093321 .
## z.diff.lag6  0.339571  0.127369   2.666 0.008598 **
## z.diff.lag7  0.055714  0.113323   0.492 0.623762
## z.diff.lag8 -0.129639  0.110829  -1.170 0.244145
## z.diff.lag9  0.134798  0.112482   1.198 0.232832
## z.diff.lag10 0.161048  0.112417   1.433 0.154251
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.01281 on 137 degrees of freedom
## Multiple R-squared:  0.2367, Adjusted R-squared:  0.1754
## F-statistic: 3.861 on 11 and 137 DF,  p-value: 7.308e-05
##
##
## Value of test-statistic is: -4.169 8.6921
##
## Critical values for test statistics:
##      1pct  5pct 10pct
## tau2 -3.46 -2.88 -2.57
## phi1  6.52  4.63  3.81

```

13) ADF-test for differences of log-transformed data (type: drift, lags: 4)

```
##
## #####
## # Augmented Dickey-Fuller Test Unit Root Test #
## #####
##
## Test regression drift
##
##
## Call:
## lm(formula = z.diff ~ z.lag.1 + 1 + z.diff.lag)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -0.085604 -0.005863  0.001323  0.007474  0.052326
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  0.004969   0.001702   2.920 0.004043 **
## z.lag.1      -0.296389   0.076771  -3.861 0.000168 ***
## z.diff.lag1   0.114928   0.088209   1.303 0.194614
## z.diff.lag2   0.131336   0.095014   1.382 0.168953
## z.diff.lag3   0.068860   0.103866   0.663 0.508373
## z.diff.lag4 -0.192558   0.103441  -1.862 0.064638 .
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.01307 on 149 degrees of freedom
```

```


## Multiple R-squared:  0.1708, Adjusted R-squared:  0.143
## F-statistic: 6.139 on 5 and 149 DF,  p-value: 3.348e-05
##
##
## Value of test-statistic is: -3.8607 7.4701
##
## Critical values for test statistics:
##      1pct  5pct 10pct
## tau2 -3.46 -2.88 -2.57
## phi1  6.52  4.63  3.81

14) KPSS-test for differences of log-transformed data (type: tau)


##
## #####
## # KPSS Unit Root Test #
## #####
##
## Test is of type: mu with 4 lags.
##
## Value of test-statistic is: 0.0682
##
## Critical value for a significance level of:
##      10pct  5pct 2.5pct  1pct
## critical values 0.347 0.463 0.574 0.739

15) Checkresiduals for for differences of log-transformed data

```



16) Fstats-test (structural change)



17) Auto.arima for differences of log-transformed data

```
## Series: myts_log_diff
## ARIMA(2,0,1)(0,0,1)[4] with non-zero mean
##
## Coefficients:
```



```
##          ar1      ar2      ma1      sma1      mean
##      -0.0946  0.8675  0.8669  -0.8603  0.0176
## s.e.   0.0505  0.0510  0.0682   0.0834  0.0012
##
## sigma^2 estimated as 0.0001485:  log likelihood=478.53
## AIC=-945.05   AICc=-944.5   BIC=-926.6
##
## Training set error measures:
##
##              ME          RMSE          MAE          MPE          MAPE          MASE
## Training set -0.0001082414 0.01199563 0.007747671 0.9198218 119.9045 0.4031201
##
##              ACF1
## Training set 0.02622321
```

18) Auto.arima for log-transformed data

```
## Series: myts_log
## ARIMA(2,0,1)(0,1,1)[4] with drift
##
## Coefficients:
##          ar1      ar2      ma1      sma1      drift
##      -0.0953  0.8679  0.8677  -0.8610  0.0044
## s.e.   0.0501  0.0507  0.0680   0.0832  0.0003
##
## sigma^2 estimated as 0.000149:  log likelihood=478.53
## AIC=-945.05   AICc=-944.5   BIC=-926.6
##
## Training set error measures:
##
##              ME          RMSE          MAE          MPE          MAPE
## Training set 3.539422e-06 0.01186853 0.007669065 -0.0003563004 0.1583222
```

```

##                               MASE          ACF1
## Training set 0.3543906 0.02770831

19) ETS model (hand-selected)

## ETS(A,A,A)
##
## Call:
## ets(y = myts_log, model = "AAA")
##
## Smoothing parameters:
##   alpha = 0.6492
##   beta  = 1e-04
##   gamma = 0.1538
##
## Initial states:
##   l = 4.4905
##   b = 0.0043
##   s = 0.043 -0.0107 -0.0012 -0.0312
##
## sigma: 0.0127
##
##      AIC      AICc      BIC
## -585.3740 -584.2051 -557.4752
##
## Training set error measures:
##                               ME      RMSE      MAE      MPE      MAPE
## Training set -0.0001195207 0.01240708 0.008217006 -0.002470129 0.1697881
##                               MASE      ACF1

```

```
## Training set 0.3797112 0.1266179
```

20) ETS model (auto-selected)

```
## ETS(M,A,M)
```

```
##
```

```
## Call:
```

```
## ets(y = myts_log, model = "ZZZ")
```

```
##
```

```
## Smoothing parameters:
```

```
## alpha = 0.7501
```

```
## beta = 1e-04
```

```
## gamma = 0.2228
```

```
##
```

```
## Initial states:
```

```
## l = 4.4984
```

```
## b = 0.0041
```

```
## s = 1.0102 0.997 1.0006 0.9922
```

```
##
```

```
## sigma: 0.0026
```

```
##
```

```
## AIC AICc BIC
```

```
## -594.6865 -593.5177 -566.7877
```

```
##
```

```
## Training set error measures:
```

```
## ME RMSE MAE MPE MAPE
```

```
## Training set 8.318883e-05 0.01230898 0.008026974 0.001650006 0.1657863
```

```
## MASE ACF1
```

```
## Training set 0.3709297 0.04725038
```

21) Forecast horizon accuracy with cross-validation

