1 Introduction

This paper aims to examine the use of historic GDP growth rate to forecast future value of GDP and then compare the forecasts to actual growth rate in the period. We use a univariate time series consisting of quarterly US GDP data to establish an AR(5)-model, which uses data from 1977-2009 to forecast. Our model end up forecasting values close to the actual growth rate and slightly higher than the long run yearly growth rate of 3% .

2 Description of data

We define the variables:

- $\Delta_4 y_t$ = D4log(GDP) $= \Delta_4 \log(GDP_t) = \log(GDP_t) - \log(GDP_{t-4})$ which is the yearly GDP growth rate
- LinFor, which is a conservative forecast assuming the linear adjustment towards the long-run growth potential 3%

Our data consists of quarterly US GDP data for the period 1975 to the second quarter of 2021. The first graph illustrates GDP, which has an upwarding trend. It implies non-stationarity in our data. The same is true for log(GDP) which is not stationary. We use the fourth difference to ensure stationarity by removing seasonal factors. When examining the $\Delta_4 y_t$ graph, it seems the trend is fluctuating around a constant mean, which implies stationarity. In addition, our data also contains the series LinFor.

The data has been downloaded as the series GDPC1 from the FRED database maintained by the Federal Reserve Bank of St. Louis.

To determine the amount of lags to include, we apply a Box-Jenkins identification method by inspecting the partial autocorrelation function (PACF) and it appears that at least two lags should be included. We also test against lags up to 10, since lag 10 is also significant in the PACF.

3 Econometric theory

3.1 AR-model

We look at an AR-model with p lags:

$$y_t = \delta + \theta_1 y_{t-1} + ... + \theta_p y_{t-p} + \epsilon_t \qquad t = 1, 2, ..., T$$

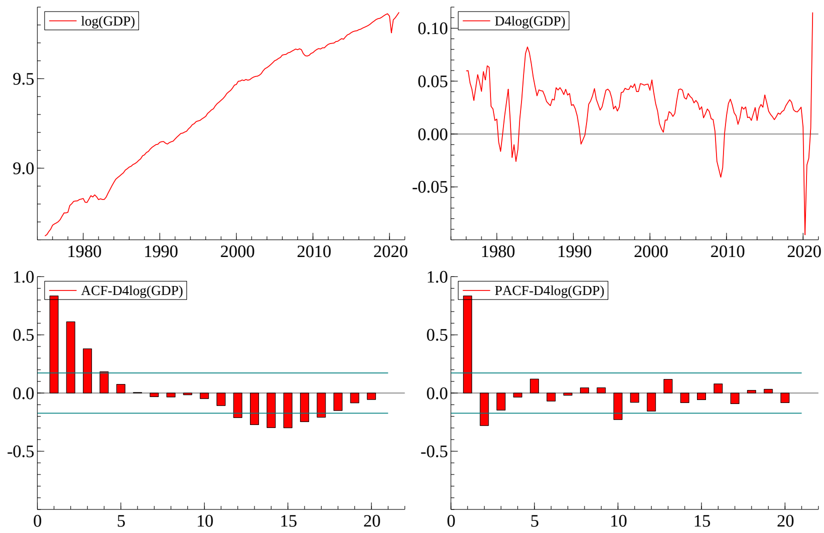

Figure 1: Data plots with ACF and PACF

y_t is the dependent variable, δ is a constant and θ_p is the coefficient for y_{t-p}. ϵ_t is the deviation from the conditional mean and is assumed to be i.i.d. $\epsilon_t \sim (0, \sigma^2)$. Furthermore $y_{t-1}, ..., y_{t-p}$ are the already observed values.

We use the Maximum likelihood estimation (MLE) to estimate the time series models as we assume a Gaussian distribution. The MLE is given as:

$$\hat{\theta}(y_1, ..., y_T) = \arg\max_{\theta \in \Theta} \log L(\theta | y_1, ..., y_T)$$

For the MLE to be consistent, the score and the Hessian contributions must obey a law of large numbers. Furthermore, the score contributions must obey a central limit theorem, and the third derivative must be bounded by a constant. We minimize the Akaike Information Criteria (AIC) to obtain the most robust model:

$$AIC = \log(\hat{\sigma}^2) + \frac{2k}{T}$$

An AR(p) model is stationary when θ_z is invertible so that the inverse root $|\phi_j| < 1$, where θ_z is the autoregressive polynomial given by:[1]

$$\theta(z) = 1 - \theta_1 z - ... - \theta_p z^p = (1 - \phi_1)(1 - \phi_2)...(1 - \phi_p)z$$

[1]Nielsen (2021) - 4.21 – 4.28

3.2 Test Misspecification

We test our model for misspecification by testing for autocorrelation, normality and heteroskedasticity problems, to secure a consistent and unbiased model.

- The test for autocorrelation is based on a Breusch-Godfrey LM test with the test statistic $\xrightarrow{d} \chi^2(h)$ under the null-hypothesis. H_0 is no autocorrelation.
- The test for heteroskedasticity is based on the White-test with the test statistic $\xrightarrow{d} \chi^2(2k)$ under the null-hypothesis. H_0 is no heteroskedasticity.
- The normality is based on a Jarque-Bera test, which uses a Goodness-of-fit to compare a normal distribution's skewness and kurtosis with that of the model. H_0 is a normal distribution, where we have the test statistic $\xrightarrow{d} \chi^2(2)$ under the null-hypothesis.

3.3 Forecast

Forecasting a univariate time series can be helpful to predict the direction of the future movements in a time series. However, it will not show turning points, and it can contain little to no economic insight. We are interested in predicting y_{T+k}.

A forecast ideally contains all the information that is observed until period T, and the forecast will be better the larger the information is. The information is defined as $\mathcal{I}_T = \{y_{-\infty}, ..., y_{T-1}, y_T\}$.

The optimal predictor of the conditional expectation is defined by:

$$y_{T+k|T} = E(y_{T+k}|\mathcal{I})$$

In a simple AR(1)-model the forecast of the next observation y_{t+1} will then be given by:

$$y_{T+1|T} = E(\delta + \theta y_T + \epsilon_{T+1}|\mathcal{I}_t) = \delta + \theta y_T$$

To judge the precision of our forecasts, we use the expected quadratic forecast error defined as:

$$c_k^2 \equiv E(y_{T+k} - y_{T+k|T})^2 = V(Y_{T+k}|\mathcal{I}_T)$$

Using this, it is possible to derive the 95% confidence intervals for the forecast of one period ahead from above:

$$y_{T+1|T} - 1.96c_1, \qquad y_{T+1|T} + 1.96c_1,$$

4 Empirical Model

4.1 Model selection

We use an AR-model to describe the univariate time series GDP, and we use a generel-to-specific method when specifying our model. According to the PACF, the lags 1, 2 and 10 seem to be statistic significant. We start by using 10 lags and then reducing our model one lag at a time. We find that an AR(5)-model has the lowest AIC, and we therefore proceed with the following model (1):

$$\Delta_4 y_t = \delta + \theta_1 \Delta_4 y_{t-1} + \theta_2 \Delta_4 y_{t-2} + \theta_3 \Delta_4 y_{t-3} + \theta_4 \Delta_4 y_{t-4} + \theta_5 \Delta_4 y_{t-5} + \epsilon_t$$

Further estimating the above model, we find lags 2 and 3 to be insignificant. Therefore, we reduce our AR model to only include lags 1, 4 and 5 (2). Our initial test results indicate heteroskedasticity, which is fixed by using robust standard errors. As normality problems are also present due to specific large residuals, we include dummy variables for those residuals and the model no longer suffer from normality problems (3).

Comparing model (3) in table 1 against model (1) and (2), we find that model (3) is has the lowest AIC, does not suffer from auto-correlation, heteroskedasticity and normality errors so it is rather robust as illustrated in figure 2.

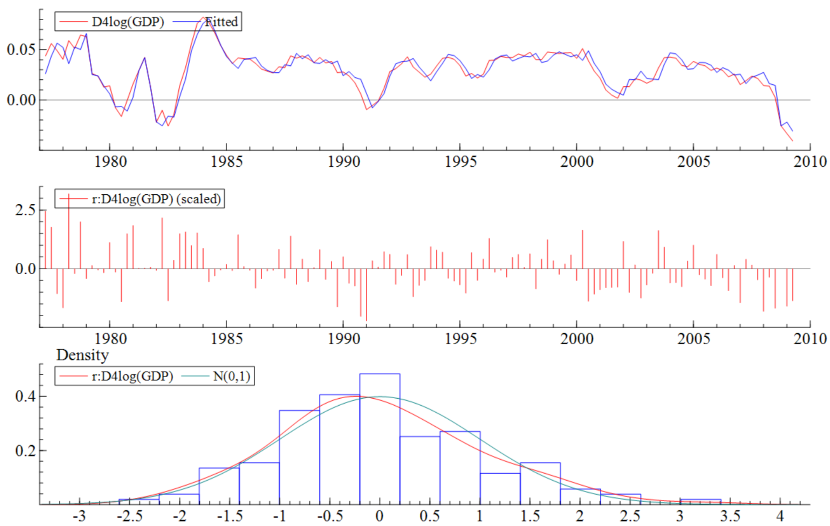

Figure 2: Actual and fitted values, scaled residuals and residual density of an AR(5) with large residual dummies

	(1)	(2)	(3)
Constant	0.003784 (2.03)	0.003907 (2.09)	0.00643 (4.34)
D4log(GDP)_1	1.179 (13.4)	1.045 (21.9)	1.036 (28.2)
D4log(GDP)_2	−0.199 (−1.49)	.	.
D4log(GDP)_3	0.0009972 (0.00731)	.	.
D4log(GDP)_4	−0.372 (−2.74)	−0.461 (−4.71)	−0.427 (−5.63)
D4log(GDP)_5	0.2516 (2.75)	0.2726 (3.01)	0.207 (2.87)
I:1980(2)+I:1979(2)	.	.	−0.02967 (−5.72)
I:1982(1)+I:1981(4)	.	.	−0.03462 (−6.46)
I:2008(4)	.	.	−0.03044 (−4.19)
$\hat{\sigma}$	0.009367	0.009419	0.007176
Log-lik.	422.532	420.778	457.435
AIC	-6.458	-6.462	-6.983
HQ	-6.404	-6.426	-6.920
SC/BIC	-6.325	-6.373	-6.828
No autocorr. 1-5	[0.22]	[0.19]	[0.30]
No hetero.	[0.05]	[0.02]	[0.23]
Normality	[0.00]	[0.00]	[0.16]
T	129	129	129
Sample start	1977(2)	1977(2)	1977(2)
Sample end	2009(2)	2009(2)	2009(2)

Table 1: Estimation results. T-ratios in $(\cdot)$ and p-values in $[\cdot]$ for misspecification tests.

Alternatively, we could have used an ARMA-model when forecasting the GDP. When we test the model (3) against an ARMA(1,1)-model with 5 lags, we end up finding that the AIC were higher in the ARMA-models compared to model (3). Therefore, we stick to the reduced AR(5)-model with dummies. Furthermore, we check for stationarity in figure 3. We see that our model is stationary, because the roots of autoregressive polynomial, $\theta(z)$, lie within the unit circle. The final model is then an AR(5) without lag 2+3 and including 5 dummies:

$$\Delta_4 y_t = \delta + \theta_1 \Delta_4 y_{t-1} + \theta_4 \Delta_4 y_{t-4} + \theta_5 \Delta_4 y_{t-5} \\ + I_{1979(2)} + I_{1980(2)} + I_{1981(4)} + I_{1982(1)} + I_{2008(4)} + \epsilon_t$$

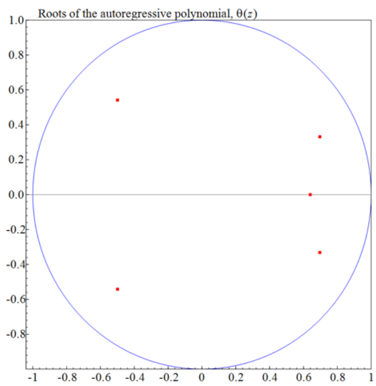

real	imag	modulus
0.6968	0.3315	0.7717
0.6968	-0.3315	0.7717
-0.4992	0.5419	0.7368
-0.4992	-0.5419	0.7368
0.6403	0.0000	0.6403

Figure 3: Unit Roots of the autoregressive polynomial

4.2 Forecasts

We use model (3) to predict the future growth of GDP after 2009(2). Figure 4 shows a forecast of GDP, a conservative forecast, LinFor, as well as the historic values of GDP.

The red line indicates the actual values of GDP, the blue line represents the forecasts and the error fans represents the 95% confidence intervals. The purple line indicates LinFor, which is a linear adjustment towards the long growth potential of 3%. We observe that the forecast of our model is higher than conservative estimate. However, the difference is rather minimal.

When we compare the forecast with the actual yearly growth in GDP from 2009 and forwards, most of the data lies beneath the forecasts, which indicates that the yearly growth in GDP is mostly lower after 2009 compared to before. Furthermore, we can also see that the actual data lies within the 95%-confidence interval, and thus our forecasts appears to be coinciding with both the actual values and LinFor.

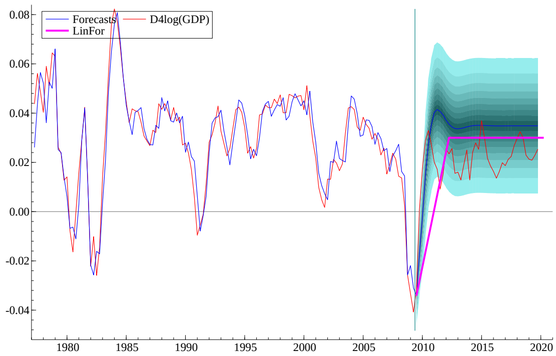

Figure 4: Forecast of GDP

5 Discussion and concluding remarks

Our series of forecasting is out-of-sample, since we forecast from 2009 to 2018, while the model is based on data from 1977 to 2009. In this case, it is not always true that the model with the lowest value of AIC is the one that provides the best out-of-sample prediction.

Overparametrization can be a problem, since the estimated model might pick up accidental patterns which has no structural meaning. However, since we only use a univariate model with lag 1, 4 and 5, overparametrization is likely not a problem. Another problem is any errors made in model selection will result in the forecast being inaccurate. Since the true data generating model is unknown, this problem is rather difficult to control for.

The last problem is that the true data generating process (DGP) may change over time which is seen in our model, since the yearly growth in GDP was larger from 1977 to 2009 compared to the period after. This might indicate that the true DGP have changed. Using an ARMA-model instead of the chosen AR-model would still face the same problems, and the problem is therefore difficult to avoid.

Overall even though our forecasts is within the 95% confidence interval of the actual values, it is not necessarily a good prediction due to the above, but it is the best forecasts, we are able produce.