Government of the
Cook Islands

Cook Islands time series economic forecasting model

Working Paper No. 18/1

1 August 2018

Prepared by the Economics Division, Ministry of Finance and Economic Management.

Table of contents

1 Summary 1

2 Introduction 3

2.1 The Cook Islands economy 3

2.2 Why do we need to forecast GDP? 5

3 Modelling approach 7

3.1 Introduction 7

3.2 Modelling principles 8

4 The ARIMA method 9

4.1 Introduction 9

4.2 The model selection process 10

5 Time series data analysis 12

5.1 Dependent variable – GDP 12

5.2 Potential explanatory variables 14

5.3 Preliminary data assessment 17

6 Model identification 19

6.1 Stationarity and cointegration 19

6.2 Preferred model specification 21

7 Parameter estimation 22

8 Diagnostic checking 23

8.1 Introduction 23

8.2 Visual methods 23

8.3 Statistical tests 25

9 Assessing forecast accuracy 26

9.1 Dynamic forecast test 26

9.2 Static forecast test 27

10 Model performance against principles 29

11 Model forecasts 30

11.1 Introduction 30

11.2 Forecast combinations 30

11.3 Conditioning the forecast explanatory variables 31

11.4 GDP forecasts 33

List of tables

Table 1-1: GDP forecasts to 2022, $m 2
Table 5-1: OLS regression summary results 18
Table 6-1: Cointegration results, Releases against Temp 20
Table 7-1: ARIMA GDP forecasting model parameters 22
Table 9-1: Measures of dynamic forecast accuracy, quarterly real GDP 27
Table 9-2: Measures of static forecast accuracy, quarterly real GDP 28
Table 10-1: Performance against principles 29
Table 11-1: GDP forecasts to 2022, $m 35

List of figures

Figure 1-1: Annual change in real GDP, percentage 2
Figure 2-1: Real annual GDP growth (percentage change) 3
Figure 2-2: Cook Islands economy by industry, 2017 4
Figure 2-3: Industry contribution to real annual GDP growth, year to June 2017 (percentage points) 5
Figure 4-1: ACF and PACF quarterly real GDP 9
Figure 5-1: Total real quarterly GDP, $m 13
Figure 5-2: Real GDP – decomposed 13
Figure 5-3: Consumer Price Index 14
Figure 5-4: Building approvals, $'000 15
Figure 5-5: Visitor arrivals by country of residence 16
Figure 5-6: Total imports, $'000 16
Figure 5-7: Total electricity generation, kWh'000 17
Figure 5-8: Pairs matrix 17
Figure 8-1: ACF and PACF plots, residuals for proposed model 23
Figure 8-2: Inverse roots for proposed model 24
Figure 8-3: Spectrum plot for preferred specification 24
Figure 9-1: Dynamic forecast accuracy performance 26
Figure 9-2: Static forecast accuracy performance 28
Figure 11-1: Holt-Winters univariate GDP forecasts 31
Figure 11-2: Holt-Winters CPI forecasts 32
Figure 11-3: Holt-Winters imports forecasts 32
Figure 11-4: Holt-Winters arrivals forecasts, monthly 33
Figure 11-5: Arrivals forecasts, annual percentage change 33
Figure 11-6: Observed and forecast real GDP, by quarter 34
Figure 11-7: Annual change in real GDP, percentage 35

List of boxes

Box 1-1: Model summary 1
Box 2-1: What is Gross Domestic Product? 3
Box 4-1: Model selection process steps 11
Box 5-1: GDP implicit price deflator 12
Box 5-2: What is the Consumer Price Index? 14
Box 6-1: Stationarity and cointegration 20

1 Summary

This technical paper sets out a time series approach to forecasting Cook Islands economic output, measured in terms of Gross Domestic Product (GDP) over the 4-year period from 2018–19 to 2021–22.

A combination forecasting method is adopted, utilising a simple average of five quarterly time series models, one multivariate – a seasonal ARIMA model – and three univariate – a seasonal ARIMA model, an exponential smoothing or error, trend, seasonal (ETS) model, a TBATS model and a seasonal naïve ARIMA model (see Box 1.1).

Box 1-1: Model summary

Multivariate model:

- seasonal ARIMA (2,1,1)(1,0,1)[4] model with the following explanatory variables, with lags shown in brackets:
  - real GDP (1)
  - consumer price Index (0, 1)
  - Import value (0,1),
  - Visitor arrivals (0,1)
  - a Fourier seasonal term (S1-4, C1-4, C2-4).

Univariate models:

- seasonal ARIMA (0,1,2)(1,1,1)[4] model
- ETS (A,N,A) exponential smoothing model
- TBATS (1,{0,0},1{<4,1>}) model
- seasonal naïve ARIMA – forecast set equal to the last observed value from the same season of the year.

Figure 1-1 shows the forecast annual change in real GDP generated from the model. A growth rate of 1.8 per cent is estimated in 2017–18, with growth forecast to rise to 2.3 per cent in 2018–19 and then fall to 1.4 per cent by the final year of the forward budget period. Average annual growth over the 4-year forward period is forecast at 1.9 per cent.

Figure 1-1: Annual change in real GDP, percentage

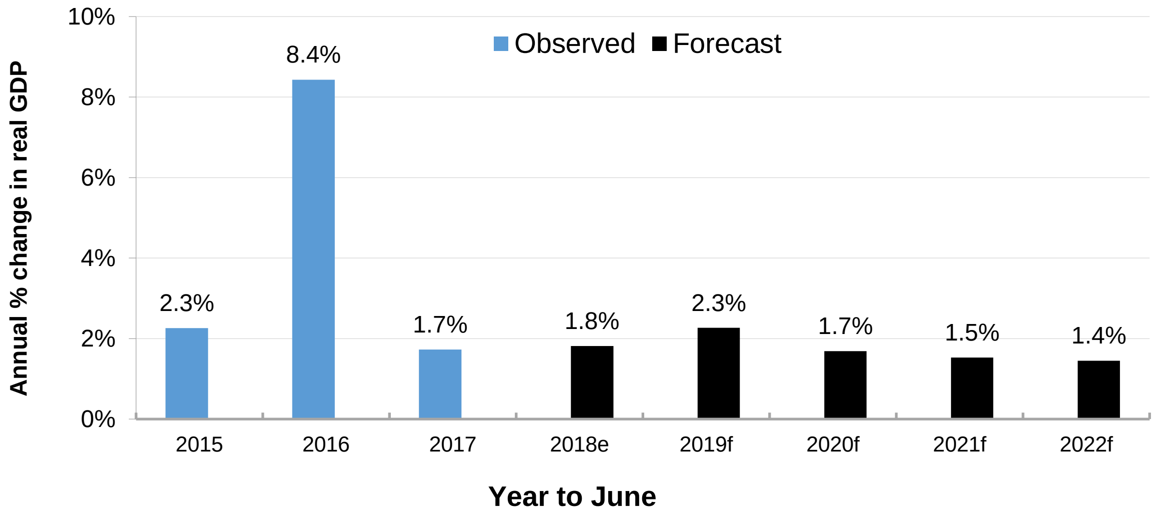

Source: MFEM analysis.

Table 1-1 provides the detailed financial year forecasts, in real and nominal terms.

Table 1-1: GDP forecasts to 2022, $m

Year to	GDP real ($m)	% change	Implicit GDP price deflator	GDP nominal ($m)	% change
Jun-14	297.4		127.77	380.0	
Jun-15	304.1	2.3%	129.18	392.9	3.4%
Jun-16	329.8	8.4%	125.87	415.1	5.6%
Jun-17	335.5	1.7%	126.18	423.3	2.0%
Jun-18	341.5	1.8%	127.16	434.3	2.6%
Jun-19	349.3	2.3%	127.62	445.7	2.6%
Jun-20	355.2	1.7%	128.31	455.7	2.2%
Jun-21	360.6	1.5%	128.83	464.6	1.9%
Jun-22	365.8	1.4%	129.45	473.5	1.9%

Source: MFEM analysis.

2 Introduction

2.1 The Cook Islands economy

The Cook Islands is a small open economy that is dominated by the tourism sector, with 164,800 tourist arrivals in the year to June 2018. Strong growth in tourism in recent years has seen the Cook Islands experience unusually high rates of economic growth.

Box 2-1: What is Gross Domestic Product?

Gross Domestic Product (GDP) is the measure of the value-added from all economic activities in the Cook Islands. Quarterly current price and constant 2006 price estimates of GDP are produced by the Cook Islands Statistics Office using the production approach. GDP estimates are disaggregated by major industry classifications, including institutional sectors.

Industry value-added is the contribution that an industry makes to the economy. A particular industry's value-added is equal to the difference in value between the goods and services inputs to the sector and the final value of the goods and services produced by the sector – or Gross Value of Production (GVP). The GVP is the total value of the goods and services produced by an industry sector. It therefore includes the value of the inputs plus the industry value-added by the sector.

Cook Islands economic output – measured in terms of real annual Gross Domestic Product (GDP) – grew by 1.7 per cent over the year to June 2017 compared to the year to June 2016, increasing from $329.7 million to $335.5 million (Figure 2-1). Although expanding at a slower rate than recent years, this continues the strong run of annual economic growth since 2014.

Figure 2-1: Real annual GDP growth (percentage change)

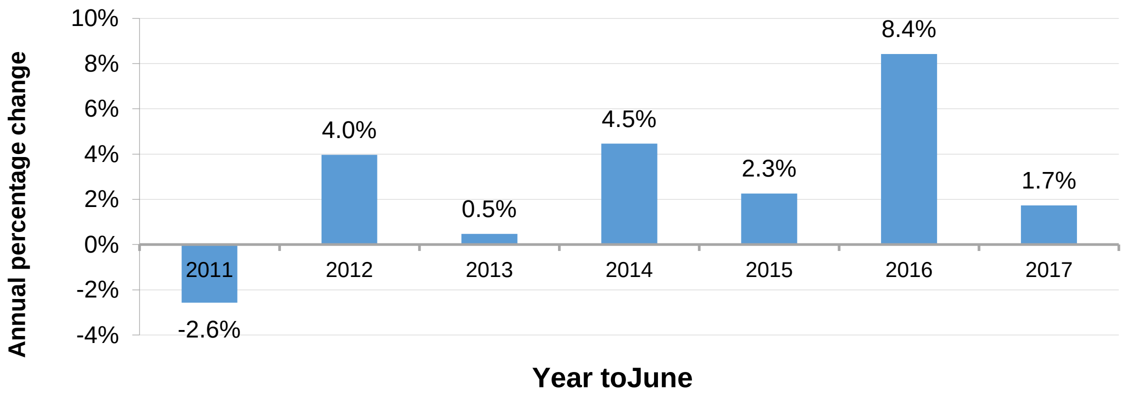

Source: Cook Islands Statistics Office, National Accounts Statistics, December Quarter 2017.

Figure 2-2 shows the breakdown of the economy by industry in 2017. The three largest industries, Restaurants & accommodation, Wholesale & retail trade and Transport & communications account for more than half of total economic output.

Figure 2-2: Cook Islands economy by industry, 2017

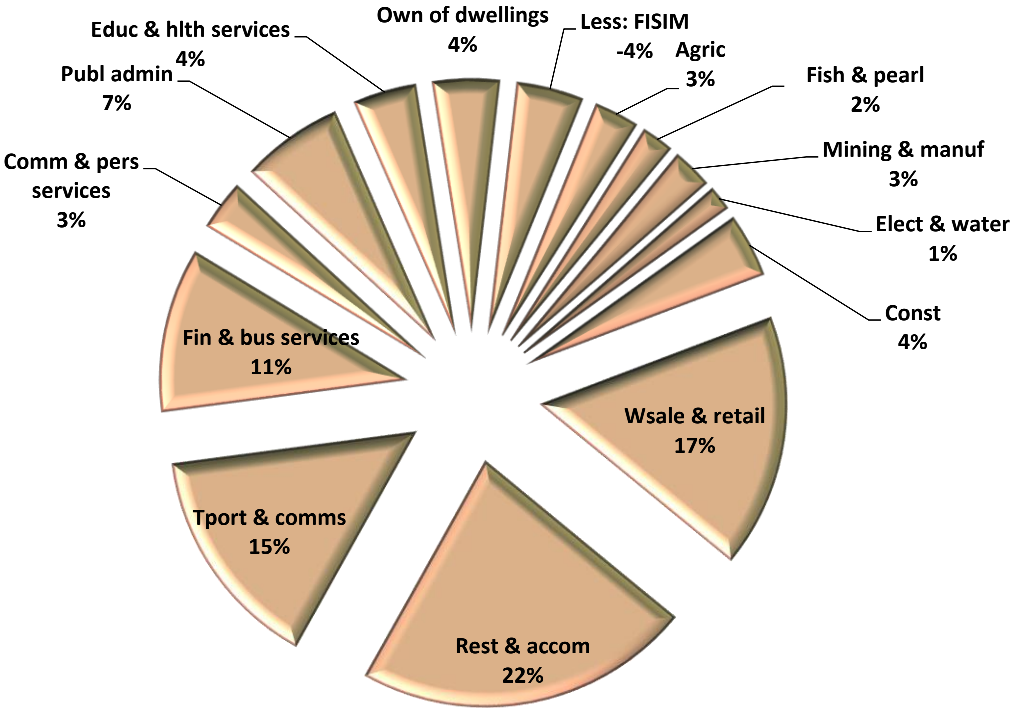

Source: Cook Islands Statistics Office, National Accounts Statistics, December Quarter 2017.

In 2016/17, high tourist arrivals drove strong growth in the economy, with the tertiary sector contributing 2.4 percentage points to economic growth. The strongest contributions were from restaurants and accommodation (1.9 percentage points), wholesale and retail trade (0.8 percentage points) and transport and communication (0.7 percentage points). Finance and business services and education and health services contracted in 2016/17. During this period, restaurants and accommodation grew by 8.4 per cent to $82 million, transport and communication grew by 4.9 per cent to $52 million and wholesale and retail trade grew by 4.6 per cent to $63 million. The secondary sector, which includes construction, electricity and manufacturing, contributed 0.5 percentage points to growth, while the primary sector, agriculture and fishing contracted, with a negative 0.5 percentage point contribution.

Figure 2-3: Industry contribution to real annual GDP growth, year to June 2017 (percentage points)

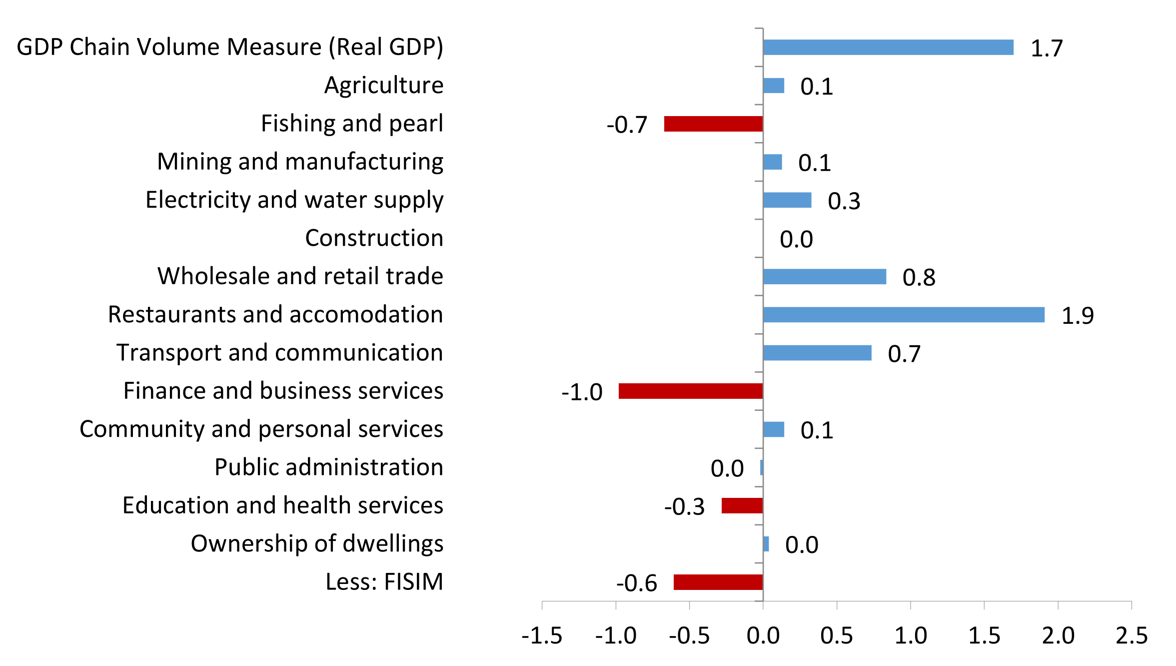

Source: Cook Islands Statistics Office, National Accounts Statistics, December Quarter 2017.

2.2 Why do we need to forecast GDP?

In the spirit of Elliot and Timmerman (2008), the overarching purpose of producing forecasts of the Cook Islands economy is to improve economic decision-making by the Cook Islands Government:

> From this perspective, forecasts do not have any intrinsic value and are only useful in so far as they help improve economic decisions.[1]

The Ministry of Finance and Economic Management (MFEM) is the central agency responsible for advising the Cook Islands Government (CIG) on financial and economic issues. The Economics Division's core role within MFEM is to provide economic advice underpinning strategic policy developments and economic opportunities in the Cook Islands. Central to this role is an understanding of current, and expected future, state of the economy and its key industries. The model presented in this paper, and its development, will contribute to this objective, as will further related work on estimating the output gap in the Cook Islands – that is the gap between actual and potential economic output.

Economic forecasts are also required on a regular basis to support the budget process, as set out in section 18 of the *Ministry of Finance and Economic Management Act 1995-96*.

The CIG annual budget estimates process determines government appropriations (operating and capital expenditure) for the upcoming financial year (1 July to 30 June), along with an estimate of government revenues. Together, this information describes the CIG's fiscal position. The budget process also provides forecasts of government expenditure and revenue

[1] Elliot and Timmerman, 2008: 4.

for an additional three years. The budget estimates process, especially the revenue side, is informed by economic growth forecasts for the 4-year period.

The current model used to forecast Cook Islands GDP is a bottom-up approach based on the expenditure national accounting measure. It uses estimates of public sector consumption and investment (including aid assistance), net merchandise trade, total tourism spending, domestic investment, and domestic consumption. Aggregate consumption expenditure is then adjusted to provide output forecasts. See section 4.6 of 2017–18 Budget Book 1 for more information on the model.[2]

The time series model presented in this paper is intended to complement the current model.

[2] http://www.mfem.gov.ck/images/2018_Cook-Islands_Budget-Book-1_Appropriation-Bill_Appropriations-and-Commentary.pdf.

3 Modelling approach

3.1 Introduction

There are a number of things to consider before in advance of the modelling process. In no particular order, these include:

- the purpose for which the forecasts are required – at the very least this will determine the forecast period, and will also inform the level of disaggregation required;
- the choice between a univariate or multivariate approach; and
- the availability and quality of data for both the dependent and, in the case of a multivariate approach, potential explanatory variables;

Each of the above matters are considered in turn below.

Purpose

As discussed in section 2.2, 4-year forecasts are required for budget purposes. The 2018-19 budget therefore requires economic forecasts for the period 1 July 2018 to 30 June 2022. The latest historical GDP data available is for the December quarter 2017.

Univariate versus multivariate

It is often the case that a simple time series model that relies solely on the relationship between previous values of the dependent variable itself will provide a better basis for forecasting than some complicated multivariate model that relies on economic theories.

> Forecasting models are best viewed as greatly simplified approximations of a far more complicated reality and need not reflect causal relations between economic variables. Indeed, simple mechanical forecasting schemes such as the random walk are often found to perform well empirically although they do not provide new economic insights into the underlying variable. …. Conversely, models aimed at uncovering true unconditional relationships in the data need not be well suited for forecasting purposes.[3]

In this paper, a combination approach is adopted, following consideration of forecasting performance in comparison to naïve alternatives. This is discussed in more detail in Chapter 9.

Data

The Cook Islands Statistics Office publishes GDP data by industry sector on a quarterly basis, along with a range of other quarterly economic data, such as a Consumer Price Index (CPI), building approvals, bank loans and bank interest rates. The Statistics office also publishes monthly visitor arrivals data, by country of origin. The data is explored in more detail in Chapter 5.

[3] Elliot and Timmerman, 2008: 4.

3.2 Modelling principles

The following modelling principles have been adopted for the purposes of this paper:

- simple, transparent and replicable
- sound statistical basis – the model should be developed using an objective model selection process that takes account of the historical relationship between the dependent and relevant explanatory variables
- fit for purpose with respect to the forecasting period – the model should be configured to produce forecasts for the period required
- perform well when tested for forecast accuracy – using dynamic out-of-sample tests in comparison to reasonable alternative models.

4 The ARIMA method

4.1 Introduction

Time series economic data are often characterised by autocorrelation (serial correlation) of the disturbances across periods, which violates the ordinary least squares (OLS) assumption that the error term is independently distributed across observations. It results in the estimates being inefficient and also adversely affects any inferences derived from the estimates.

Evidence of autocorrelation can be found by examining the autocorrelation function (ACF) which measures the linear relationship between lagged values of a time series and shows the degree of persistence over respective lags of a variable. The ACF for real GDP in Figure 4-1 shows a persistent pattern with clear evidence of autocorrelation.[4]

Figure 4-1: ACF and PACF quarterly real GDP

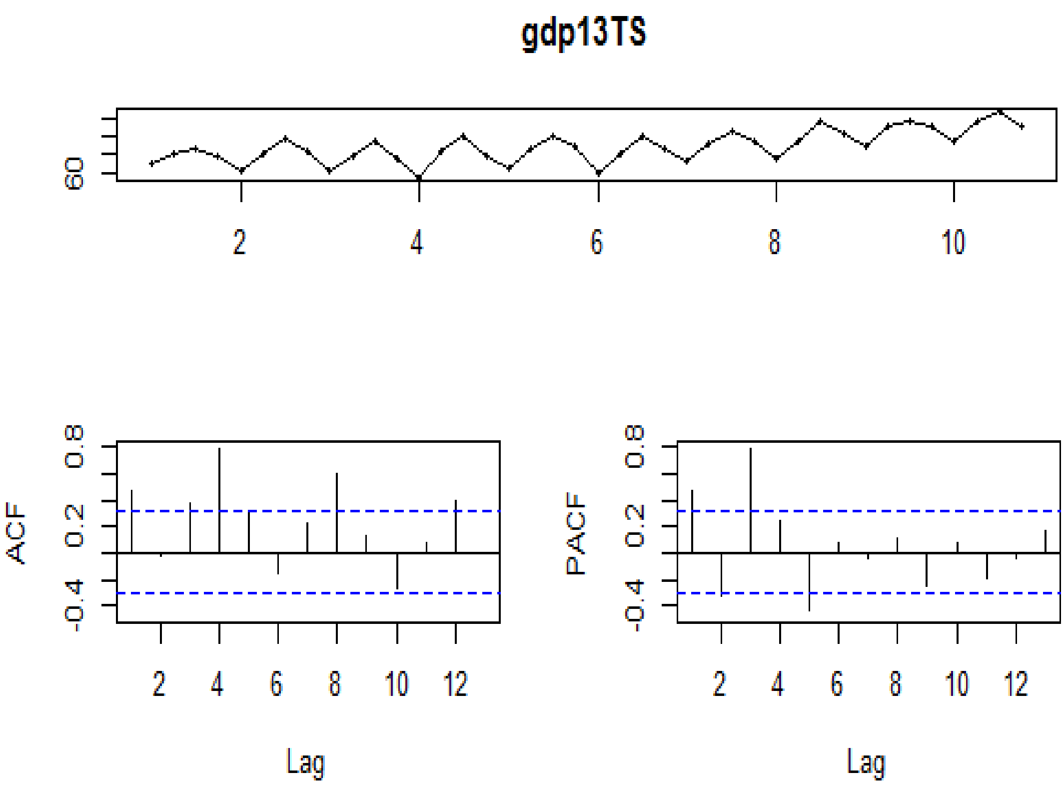

Source: MFEM analysis, R Studio output.

An examination of the partial autocorrelation function (PACF), also shown in Figure 4-1, which measures the relationship between a variable and a lag of itself after removing the effect of other time lags, shows a strong partial coefficient at lags one and five.[5]

[4] Real GDP is nominal GDP adjusted for the effect of inflation.

[5] Following (Greene 2012) and (Hyndman and Khandakar 2008), the ACF(*k*), which gives the gross correlation between y_t and y_{t-k} for different values of *k*, can mask a completely different underlying relationship. For example, a correlation between y_t and y_{t-2}, could arise simply because both variables are correlated with y_{t-1} rather than any new information contained in y_{t-2} that could be used forecasting y_t. To overcome this problem, the PACF can

The autoregressive integrated moving average (ARIMA) procedure is a time series technique designed to model the lagged relationships in serially-correlated data. The ARIMA procedure models a time series event as a linear function of its past values, past errors and current and past values and errors of other time series (in a multivariate ARIMA model).

More specifically, an ARIMA model integrates an autoregressive model (which uses past values of the forecast variable) with a moving average model (which uses past forecast errors):[6]

- autoregressive – a model that uses the dependent relationship between an observation and some number of lagged observations.
- integrated – the use of differencing of raw observations (e.g. subtracting an observation from an observation at the previous time step) in order to make the time series stationary.
- moving average – a model that uses the dependency between an observation and a residual error from a moving average model applied to lagged observations.

Following Hyndman and Athanasopoulos (2012), this can be written as:

$$y_t = c + \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + \theta_1 e_{t-1} + \cdots + \theta_q e_{t-q} + e_t$$

where y_t is the differenced series. This is called an ARIMA (p,d,q) model. ARIMA models can model non-seasonal and seasonal data. The seasonal model can be written as:

$$\text{ARIMA} \quad \underbrace{(p,d,q)}_{Non-seasonal} \; \underbrace{(P,D,Q)}_{Seasonal} m$$

where:

- *p* and *P* is the order of the autoregressive part
- *d and D* is number of differences needed for stationarity[7]
- *q* and Q is the order of the moving average part – the number of lagged forecast errors
- *m* is the number of periods per season.

The appropriate values for the parameters above are identified in the modelling process.

4.2 The model selection process

In developing a preferred ARIMA model, MFEM has followed a process (see Box 5), based on the Box-Jenkins approach (Box and Jenkins 1970)[8].

be used which measures the relationship between y_t and y_{t-2} net of the intervening effect of y_{t-1} (Greene, 2012: 757; Hyndman and Khandakar, 2008: 8.5 Non-seasonal ARIMA models).

[6] For a full treatment of the ARIMA modelling theory and practical application see Greene (2012) and Hyndman and Athanasopoulos (2012).

[7] See section 2.2.6 for more detail on stationarity.

[8] The Box-Jenkins approach to modelling ARIMA processes was described in a seminal book by statisticians George Box and Gwilym Jenkins in 1970.

Box 4-1: Model selection process steps

1. **Data analysis**:
    - select the desired dependent variable and determine the relevant explanatory variables to be tested in the model identification stage
    - plot the data to look for patterns, such as seasonality or trends in the data over time
    - assess consistency in the relationship between the dependent and explanatory variables over time to inform the choice of model estimation period.
2. **Model identification**:
    - check for stationarity and evidence of cointegration between variables (when dealing with multivariate models), then differencing the data if necessary
    - identify the potential model structure by comparing the empirical autocorrelation patterns with theoretical ones using the ACF and PACF
    - run multiple alternative model specifications and selecting the preferred specification with reference to the Akaike Information Criterion.
3. **Parameter estimation** – this step involves estimating the values of the parameters of the preferred model specification over the selected estimation period.
4. **Diagnostic checking** – the fourth stage involves examining the assumptions of the model by testing the model residuals for stationarity through visual inspection and statistical methods.
5. **Accuracy assessment** – assess forecast accuracy using a range of measures such as the root mean square error (RMSE).[9]
6. **Assessment against principles** – performance against the set of forecast approach principles set out in section 2.2.2.
7. **Forecasting** – equipped with the preferred model that has been identified, estimated and checked, the final step is to use it to compute forecasts.

Source: Adapted from Hyndman (2001).

MFEM used the following open source software in the development and operation of the ARIMA model:

- the R statistical package[10], and
- the RStudio interface.[11]

The remainder of this paper steps through the Box-Jenkins approach model selection process.

[9] The RMSE is a standard measure of the difference between the values forecast by a model and the observed values.

[10] R Core Team (2018). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/. R version 3.4.4 "Someone to Lean On" released on 15 March 2018.

[11] RStudio Team (2016). RStudio: Integrated Development for R. RStudio, Inc., Boston, MA URL http://www.rstudio.com/.

5 Time series data analysis

5.1 Dependent variable – GDP

5.1.1 Current and constant price GDP

The Australian Bureau of Statistics (ABS), in an article entitled 'Demystifying Chain Volume Measures', states that when comparing the difference in aggregate values between two time periods, any observed movement is generally a combination of changes in quantity and changes in price. Of interest to many users of economic data is understanding the degree to which the dollar value of economic growth (either positive or negative) between two periods is being driven by changes in quantities (that is the physical volumes of production and consumption) as distinct from changes in prices.

The requirement for a measure of economic growth due only to changes in quantities has resulted in the development of two types of data series in which the effects of price changes are removed: constant price estimates and chain volume measures. Both measures indicate changes in quantity (or volume) between time periods by keeping the prices of goods and services constant.

The Cook Islands Statistics Office publishes GDP uses the former measure, publishing quarterly GDP data in both current and constant prices. The 2006 calendar year is used as the base year, with current prices adjusted using an implicit price deflator (see Box 5.1).

Box 5-1: GDP implicit price deflator

Implicit price deflators for expenditure on GDP provide a broad measure of price change for total economic activity. They provide an estimate of price change between the base period and any other period, using the quantity weights in the latter period. Because weights change from period to period, a change in an implicit price deflator between any two periods reflects changes in both actual prices and weights or compositional changes. It is calculated as GDP at current prices divided by GDP at constant prices, multiplied by 100.

Source: Cook Islands Statistical Office.

For the purposes of this paper, it is the constant or 'real' price GDP that is forecast.

5.1.2 Real GDP time series

Real quarterly GDP from December 2007 to December 2017, the latest available, is shown in Figure 5-1. There is clear indication of a seasonal pattern and an increasing trend in the data,

Figure 5-1: Total real quarterly GDP, $m

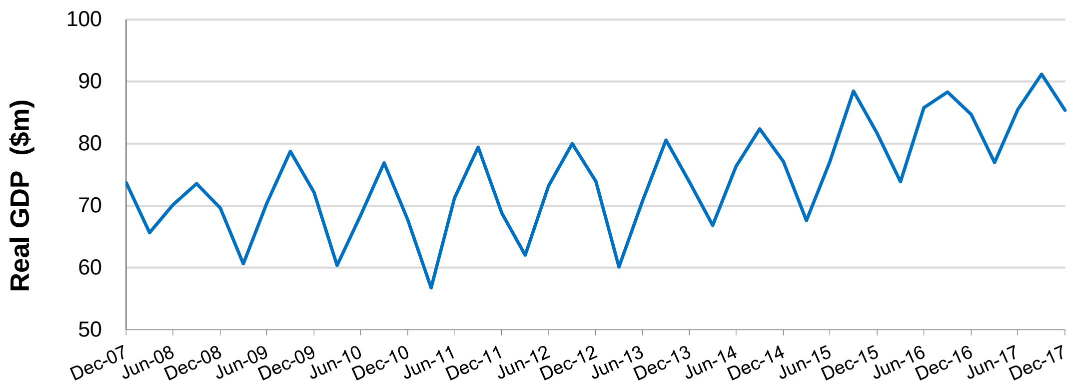

Source: Cook Islands Statistics Office

Decomposing the time series GDP data into its trend, seasonal and random components confirms the seasonal and trend pattern (see Figure 5-2).

Figure 5-2: Real GDP – decomposed

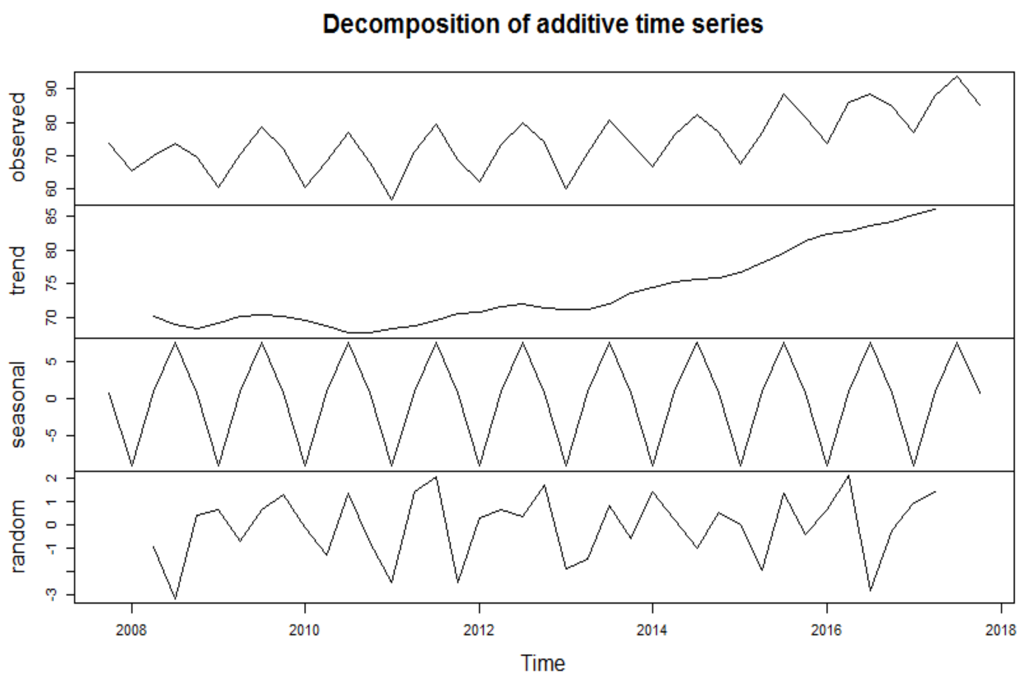

Source: MFEM analysis, R Studio output.

The seasonal pattern reflects the tourist season – peak tourist season is from May to September which matches GDP peaks in the June and September quarters, and troughs in December and March.

5.2 Potential explanatory variables

5.2.1 Introduction

A number of potential explanatory variables were assessed for the purposes of developing a multivariate model. The final choice was limited to data available in quarterly (or monthly) frequency over the 10-year period from December 2007 to match the quarterly GDP time series. For example, one key potential explanatory variable, the unemployment rate, was excluded on this basis.

5.2.2 Consumer Price Index

The Cook Islands Statistical Office publishes a Consumer Price Index (CPI) each quarter that measures changes in prices on Rarotonga (see Box 5.1 for more detail).

Box 5-2: What is the Consumer Price Index?

The Consumer Price Index covers price changes of a basket of goods and services consumed by all households on Rarotonga (the main island).

The base year is 2006. Prices are collected for 205 items and from selected outlets around Rarotonga. Individual prices are combined using weights from the Household Expenditure Survey (HES) conducted in 2004/5. The HES information was used to select the basket of goods and services.

The CPI is computed using the Laspeyres price index formula.

Source: Cook Islands Statistical Office.

The CPI for the period December 2007 to March December 2017 is shown in Figure 5-3. The CPI data shows an inclining trend from December 2007 until June 2015, followed by a declining trend from then until December 2017.

Figure 5-3: Consumer Price Index

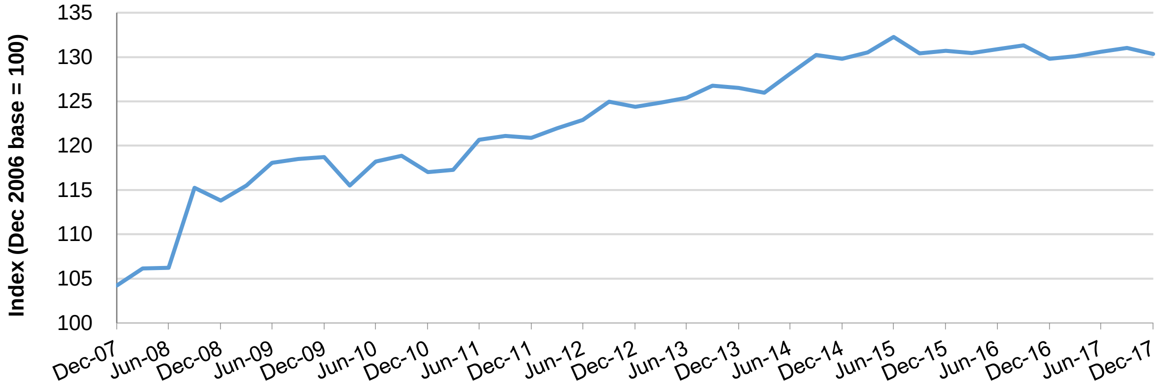

Source: Cook Islands Statistics Office.

5.2.3 Building approvals

The Statistics Office publishes quarterly data on the value of building approvals in the Cook Islands, disaggregated into residential, commercial and community categories, as shown in Figure 5-4.

Figure 5-4: Building approvals, $'000

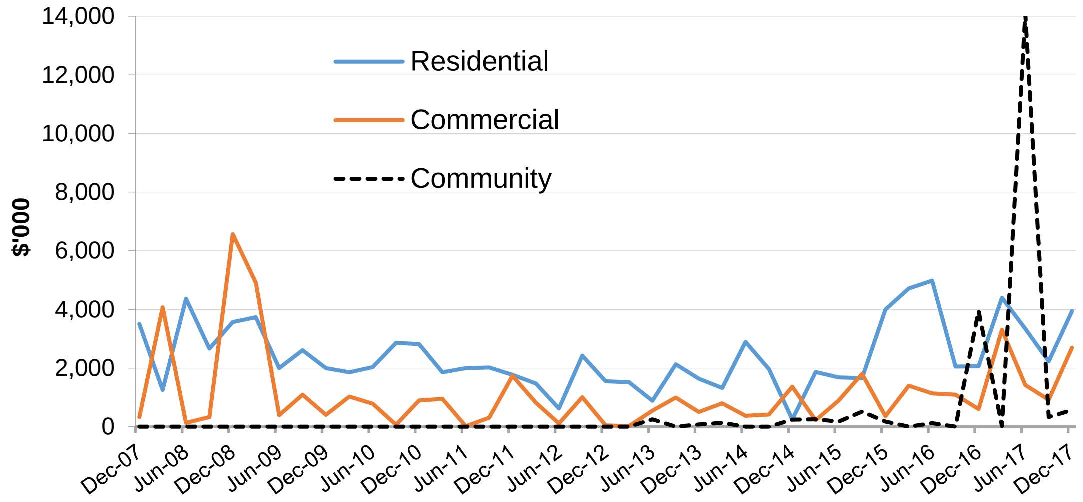

Source: Cook Islands Statistics Office.

5.2.4 Visitor arrivals

The Statistics Office publishes monthly data on visitor arrivals in the Cook Islands, disaggregated by country of usual residence. Figure 5-5 shows monthly arrivals data from January 1995 to May 2018. The data shows a strong inclining trend and a very clear seasonal pattern, with peaks in the cooler 'dry' months from May to September, and troughs in the warmer and wetter period from October to April.

Figure 5-5: Visitor arrivals by country of residence

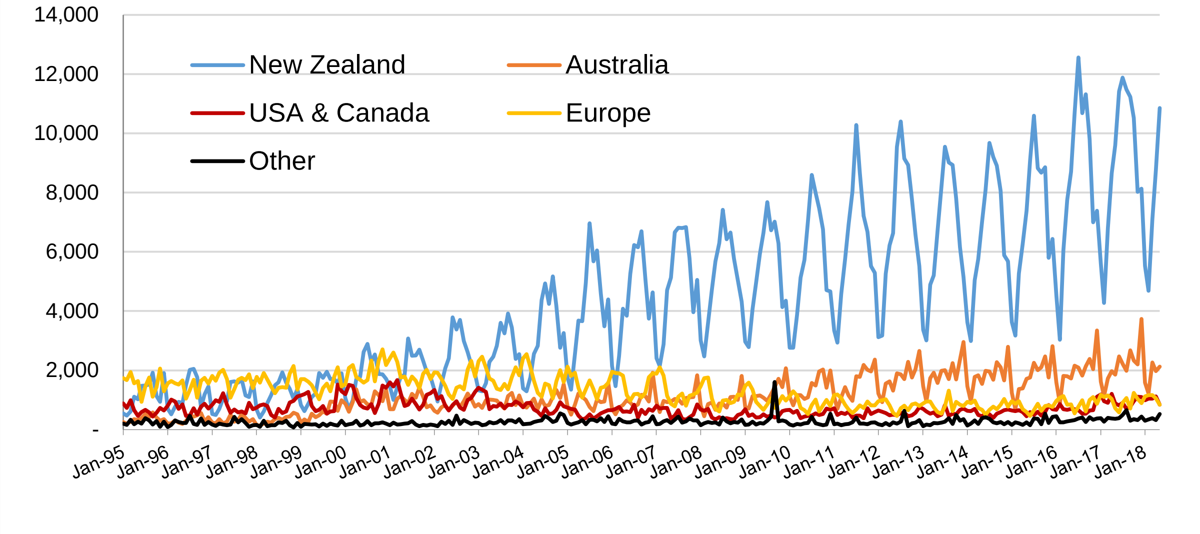

Source: Cook Islands Statistics Office.

5.2.5 Imports

The Statistics Office publishes quarterly data on the value of goods and services imported into the Cook Islands, as shown in Figure 5-6Figure 5-4, disaggregated by industry category and country of origin. The data shows a declining trend over the first five years, followed by an inclining trend, and a fairly regular seasonal pattern, with peaks in the September quarter.

Figure 5-6: Total imports, $'000

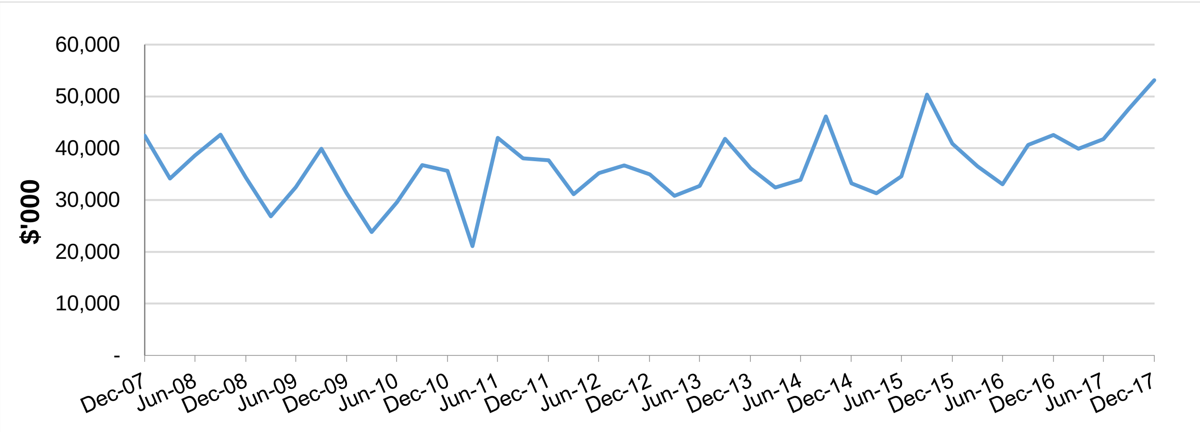

Source: Cook Islands Statistics Office.

5.2.6 Electricity generation

The Statistics Office publishes quarterly data on electricity generation in the Cook Islands, as shown in Figure 5-7, disaggregated by island. The data shows no trend and no obvious seasonal pattern.

Figure 5-7: Total electricity generation, kWh'000

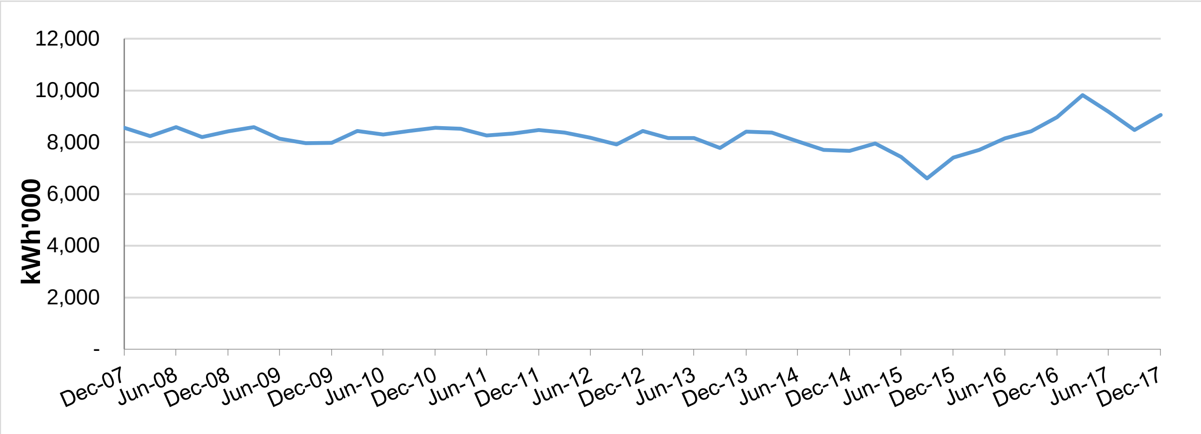

Source: Cook Islands Statistics Office.

5.3 Preliminary data assessment

The 'pairs' function in R produces a matrix of scatter plots that provides a visual indication of the correlation between the dependent variable (RGDP) and potential explanatory variables. Three of the potential variables discussed in this chapter – CPI, imports (IMP) and arrivals (ARR) – show a positive correlation with RGDP. Of the three, CPI appears to show the weakest relationship with RGDP.

Figure 5-8: Pairs matrix

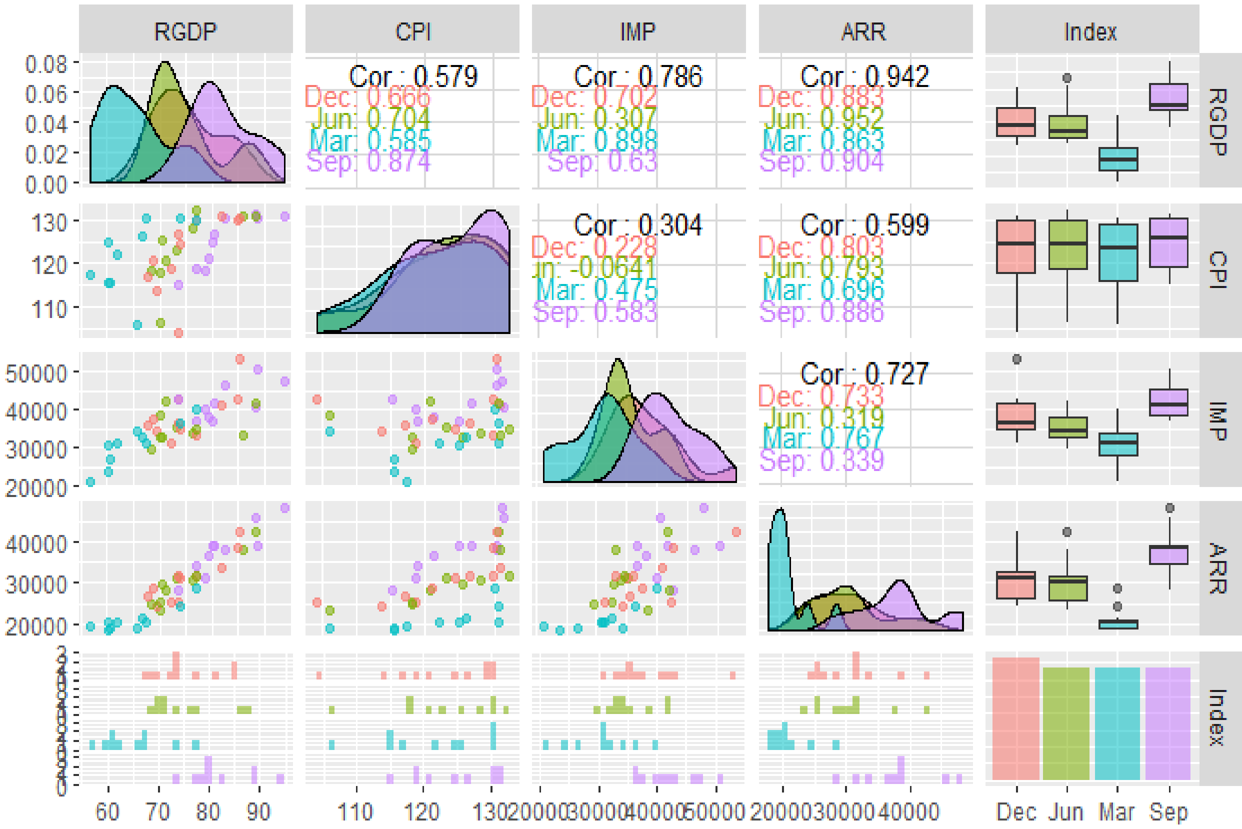

Source: MFEM analysis.

The results of an OLS regression of RGDP against the three identified variables is shown in Table 5-1. In line with the correlation result above, the IMP and ARR coefficients are significant, with a multiple R-squared of 0.91. The F-statistic is also significant at the 95 per cent level.

Table 5-1: OLS regression summary results

	Estimate	Standard error	P-value	Significance
Intercept	28.76	8.37	0.00	**
CPI	0.08	0.07	0.24	
IMP	0.00	0.00	0.00	**
ARR	0.00	0.00	0.00	***

Source: MFEM analysis.

6 Model identification

6.1 Stationarity and cointegration

The methods used to estimate ARIMA models rely on the residuals or error term being serially uncorrelated (or white noise). A stationary process exhibits a constant mean and variance over time:

> A stationary time series is one whose properties do not depend on the time at which the series is observed.[12]

What this means is that a seasonal time series such as Cook Islands GDP is not stationary by definition. Differencing, that is calculating the difference between consecutive observations, is a common technique to stabilise a non-stationary series.

> Accumulated wisdom and the results of the previous sections suggest that the appropriate way to manipulate such series is to use differencing and other transformations (such as seasonal adjustment) to reduce them to stationarity and then to analyze the resulting series as VARs or with the methods of Box and Jenkins.[13]

However, in certain circumstances, the error process of a multivariate ARIMA model may be stationary despite evidence of non-stationarity in the individual variables. This occurs in cases where pairs of variables are cointegrated see Box 6.1). In this particular circumstance, it is inappropriate to difference the variables as the 'estimated coefficients [of the equation] are correct'.[14]

This issue was explored by first assessing the need for differencing and second by applying a statistical test for cointegration.

The 'nsdiffs' function in the R 'forecast' package was used to estimate the number of differences required to make the GDP and three potential explanatory variable time series data stationary.[15] The results were as follows:

- all four series require differencing once; and
- RGDP, IMP and ARR require seasonal differencing once;

[12] Hyndman and Athanasopoulos, 2012: section 8.1 Stationarity and differencing.

[13] Greene, 2012: 999.

[14], Hyndman and Athanasopoulos, 2012: section 9.1 Dynamic regression models.

[15] Hyndman and Khandakar, 2008: 1-22.

Box 6-1: Stationarity and cointegration

In the fully specified regression model

$$y_t = \beta x_t + \varepsilon_t$$

there is a presumption that the disturbances ε_t are a stationary, white noise series.[16]

But this presumption is unlikely to be true if y_t and x_t are integrated series. Generally, if two series are integrated to different orders, then linear combinations of them will be integrated to the higher of the two orders. Thus, if y_t and x_t are I (1)—that is, if both are trending variables—then we would normally expect $y_t - \beta x_t$ to be I (1) regardless of the value of β, not I (0) (i.e., not stationary). If y_t and x_t are each drifting upward with their own trend, then unless there is some relationship between those trends, the difference between them should also be growing, with yet another trend. There must be some kind of inconsistency in the model. On the other hand, if the two series are both I (1), then there may be a β such that

$$\varepsilon_t = y_t - \beta x_t$$

is I (0). Intuitively, if the two series are both I (1), then this partial difference between them might be stable around a fixed mean. The implication would be that the series are drifting together at roughly the same rate. Two series that satisfy this requirement are said to be cointegrated, and the vector [1,−β] (or any multiple of it) is a cointegrating vector.

Source: Greene (2012).

As the results suggest that all variables considered are non-stationary of order 1, the next step was to test for cointegration. The 'urca' package was used to apply the Phillips and Ouliaris unit root test to check for evidence of cointegration between the Releases and Temp data (Pfaff 2008).[17] As shown in Table 2-5, the null hypothesis of no cointegration cannot be rejected at the 95 per cent level for any combination tested. This result indicates that differencing should be applied in the model estimation.

Table 6-1: Cointegration results, Releases against Temp

	Test statistic	10% critical value	5% critical value	1% critical value
RGDP-CPI	23.174	20.39	25.97	38.34
RGDP–IMP	2.890	20.39	25.97	38.34
RGDP–ARR	0.832	20.39	25.97	38.34

Source: MFEM analysis.

[16] If there is autocorrelation in the model, then it has been removed through an appropriate transformation.

[17] For a fuller treatment of the test see Phillips and Ouliaris, 1990: 165-193.

6.2 Preferred model specification

6.2.1 Process

The next step in the model identification process involved running multiple model specifications in an iterative fashion. The preferred specification was identified on the basis of minimising the Akaike Information Criterion (AIC), significance of the parameter estimates, stationarity of the model residuals and unit root performance (see Chapter 8) and forecast accuracy (see Chapter 9). The AIC is a statistical measure that values model fit and parsimony, where, all else being equal, the minimum value of the AIC is generally the best model for forecasting:

$$AIC = N\ log\left(\frac{SSE}{N}\right) + 2(k + 2)$$

where:

- *N* is the number of observations in the estimation
- *k* is the number of predictors in the model
- *SSE* is the minimum sum of squared errors.

The preferred model was identified and estimated using time series data from December 2007 to December 2018.

6.2.2 Preferred model

The preferred GDP forecasting model is a seasonal ARIMA (2,1,1)(1,0,1)[4] model with the following explanatory variables, with lags shown in brackets:

- RGDP (1)
- CPI (0, 1)
- IMP (0,1),
- ARR (0,1)
- a Fourier seasonal term (S1-4, C1-4, C2-4).

The Fourier series was included to capture any seasonality not captured by the quarterly seasonality in the ARIMA process, or explained by the explanatory variables that display a seasonal pattern. Following Hyndman (2010), the Fourier series approach allows the seasonal pattern to be modelled as follows:

$$\mathrm{y_t} = \mathrm{a} + \sum_{\mathrm{k}=1}^{\mathrm{K}}\left[\alpha \sin\left(\frac{2\pi \mathrm{kt}}{\mathrm{m}}\right) + \ \beta \sin\left(\frac{2\pi \mathrm{kt}}{\mathrm{m}}\right)\right] + \mathrm{N_t}$$

where:

- $\mathrm{N_t}$ summarises all the other variables in the model, including the ARIMA error term.

A value of *K* of 2 was found to minimise the AIC.

7 Parameter estimation

The parameters of the preferred model are shown in Table 7-1. The model returns an AIC of 164.74 with the majority of the explanatory variables significant at the 99 per cent level.

Table 7-1: ARIMA GDP forecasting model parameters

Variables	Coefficient	p-value	Sig
ar1	-0.01	0.89	not sig
ar2	-0.86	0.00	***
ma1	-1.00	0.00	***
sar1	-0.36	0.04	**
sma1	-1.00	0.00	***
rgpdp1	0.83	0.00	***
cpi0	-0.17	0.06	*
cpi1	0.31	0.00	***
imp0	0.00	0.00	***
imp1	0.00	0.00	***
arr0	0.00	0.00	***
arr1	-0.00	0.00	***
S1-4	-4.53	0.00	***
C1-4	-4.27	0.00	***
C2-4	0.81	0.00	***

Source: MFEM analysis, R Studio output.

8 Diagnostic checking

8.1 Introduction

As noted earlier, the methods used to estimate ARIMA models rely on the model residuals or error term being serially uncorrelated (or white noise). After identifying and estimating the preferred model, the fourth step in the model selection process is an examination of the residuals of the model to check that they have no remaining autocorrelations. This involves visual inspection of the ACF, unit root and spectrum charts, and statistical tests on the residuals.

8.2 Visual methods

Figure 8-1 shows the ACF and PACF of the model residuals, which appear to be white noise as all the spikes fall within the significance limits (shown by the dotted blue lines).

Figure 8-1: ACF and PACF plots, residuals for proposed model

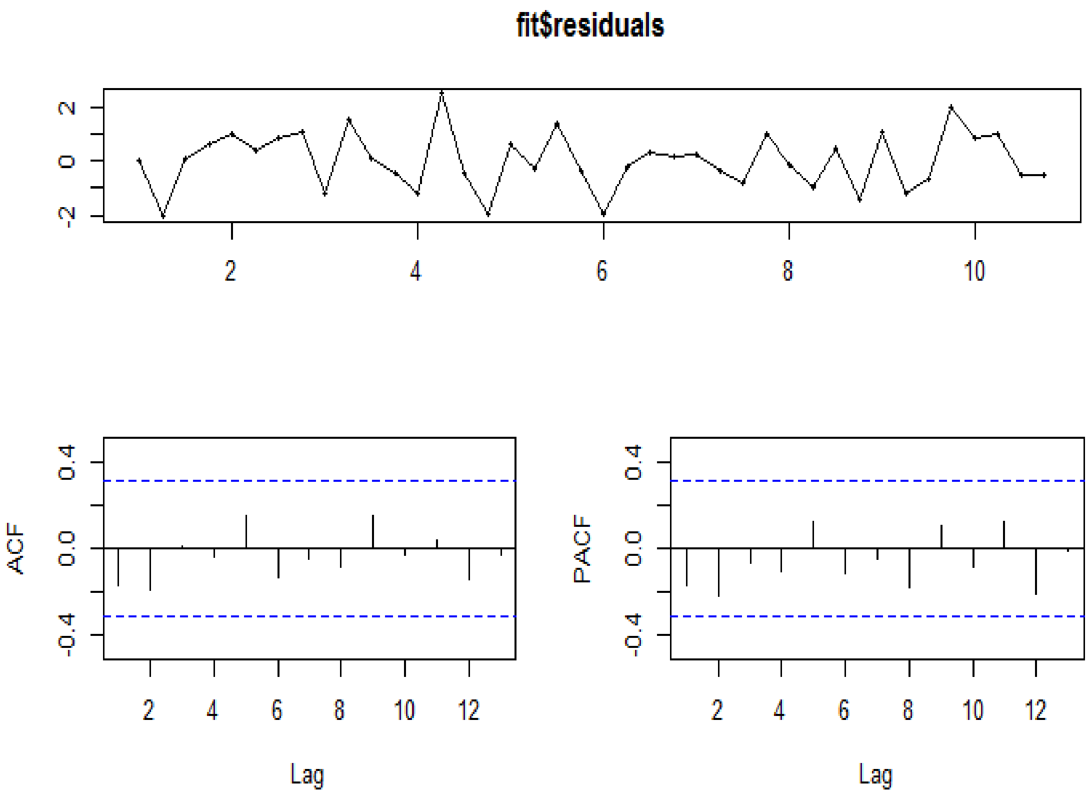

Source: MFEM analysis, RStudio output.

The stationarity of an ARIMA process depends on the AR parameters. If the inverse roots of the AR polynomial all lie within the unit circle, as shown in Figure 8-1, the inference is that the estimated process is stationary.

Figure 8-2: Inverse roots for proposed model

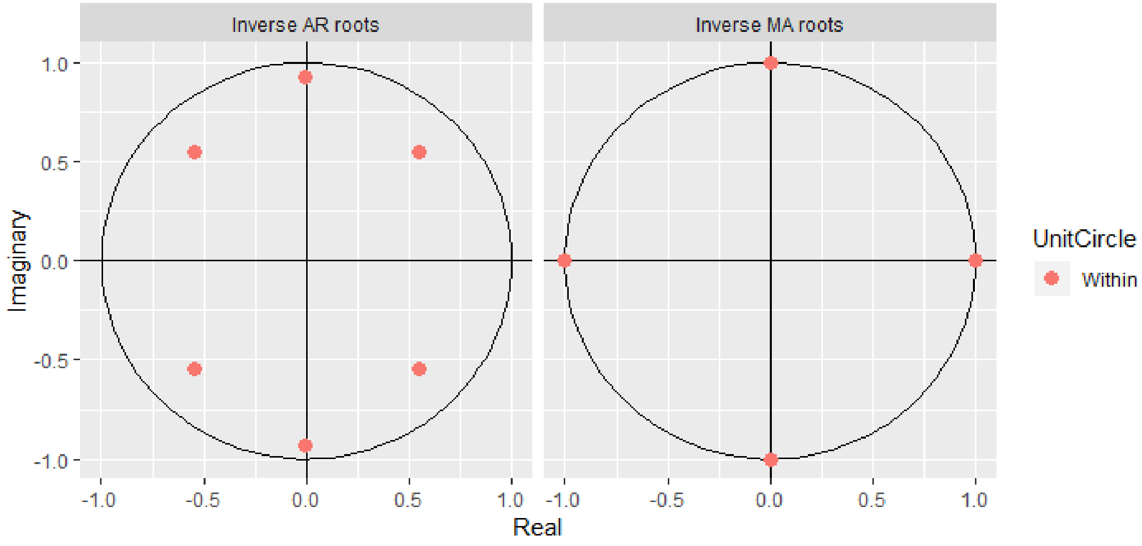

Source: MFEM analysis, RStudio output.

In addition, as shown in Figure 8-3, the spectral density of the ARIMA model residuals is reflective of a Gaussian white noise signal.[18] This is further indication that there is no remaining residual autocorrelation.

Figure 8-3: Spectrum plot for preferred specification

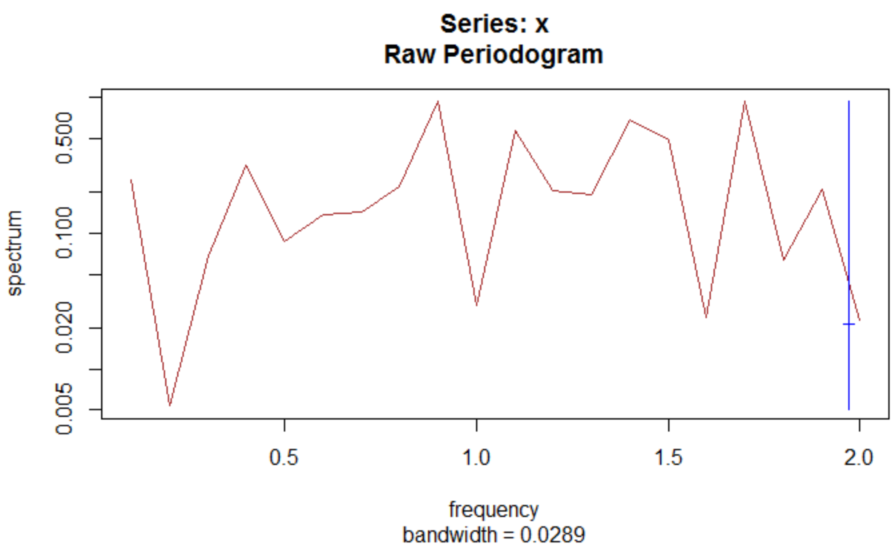

Source: MFEM analysis, RStudio output.

[18] Each element in a white noise time series is a random draw from a population with zero mean and constant variance (Greene, 2012: 749).

8.3 Statistical tests

The Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test is commonly used to check that the residuals from a time series model resemble white noise.[19] The KPSS test returned a test statistic of 0.0789, smaller than the 10 per cent significance critical value of 0.347. This suggests that the null hypothesis of white noise errors should be accepted at any reasonable significance level.

[19] See Kwiatowski, Phillips, Schmidt and Shin, 1992: 159-178, for more detail on the KPSS test.

9 Assessing forecast accuracy

9.1 Dynamic forecast test

Dynamic out-of-sample forecast performance, where in-sample model parameters are used to generate forecasts conditioned on the observed out-of-sample explanatory variables, is generally considered a good test of an econometric model.[20] Hyndman (2014) states:

> It is important to evaluate forecast accuracy using genuine forecasts. That is, it is invalid to look at how well a model fits the historical data; the accuracy of forecasts can only be determined by considering how well a model performs on new data that were not used when estimating the model.[21]

The dynamic forecast test requires withholding a portion of the sample data – the test data – from the estimation and using the rest of the data – the training data – for estimating the model. In this case, the 3 year period from March 2015 was withheld as the test data, with the model estimated using training data from December 2007 to December 2014.

9.1.1 Visual inspection

The performance between the predicted dynamic values and observed releases is shown in Figure 9-1, which suggests a reasonably close relationship. At the aggregate level, the dynamic forecast predicts total GDP of $688 million for the test period, about 2 per cent above the observed $672 million.

Figure 9-1: Dynamic forecast accuracy performance

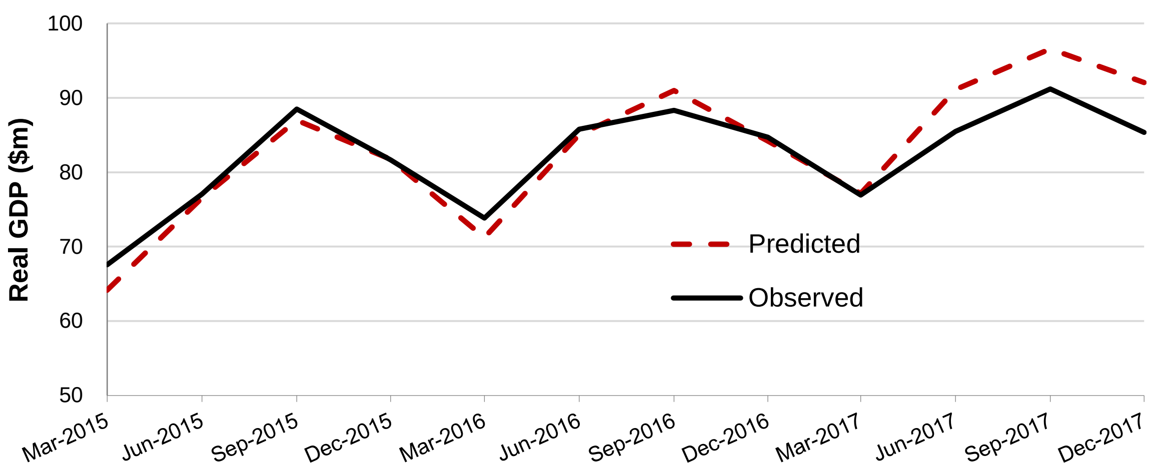

Source: MFEM analysis.

[20] Clements and Hendry, 2003: 1.

[21] Hyndman, 2014: 1.

9.1.2 Forecast accuracy measures

As the model is a quarterly GDP forecasting model, quarterly performance is examined, in comparison with alternative forecasting models. The seasonal naïve ARIMA (snaïve) model and exponential smoothing (ETS) models have been selected for this purpose.

There are a range of common measures of forecast accuracy, four of which are applied here:

- mean error: $ME = mean(e_i)$
- root mean squared error: $RMSE = \sqrt{mean(e_i^2)}$
- mean absolute percentage error: $MAPE = mean(100e_i/y_i)$
- Theil's U1 statistic:[22] $U1 = \sqrt{\sum_{i=1}^{n}(y_t - \hat{y}_t)^2} \Big/ \left(\sqrt{\sum_{i=1}^{n} y_i^2} + \sqrt{\sum_{i=1}^{n} y_i^2}\right)$

where:

- $e_i = y_i - \hat{y}_i$
- y_i is the *i*th observation
- $\hat{y}_i$ is a forecast of y_i.

Table 9-1 shows the range of measures for the forecasts from the multivariate and alternative models, both derived from the quarterly forecast performance over the test period. The results suggest that, on almost all measures presented, the multivariate GDP model outperforms the alternatives in terms of forecast accuracy.

Table 9-1: Measures of dynamic forecast accuracy, quarterly real GDP

Measures	Multivariate	Seasonal naïve	ETS
Mean error	0.95	2.55	2.14
Root mean squared error	3.35	4.07	3.71
Mean absolute percentage error	3.02	3.42	3.41
Theil's U1	0.02	0.01	0.01

Source: MFEM analysis.

9.2 Static forecast test

The results of a static out-of-sample forecast performance for the test period for the multivariate model, with the in-sample model parameters used to generate forecasts conditioned on the observed out-of-sample dependent and explanatory variables, is shown Figure 2-20. The static forecast performance, shown in Table 9-2 demonstrates relatively low levels of variance which indicates that the preferred ARIMA model specification remains stable over the out-of-sample period.

[22] The closer Theil's U1 measure is to zero, the greater the forecasting accuracy of the model.

Figure 9-2: Static forecast accuracy performance

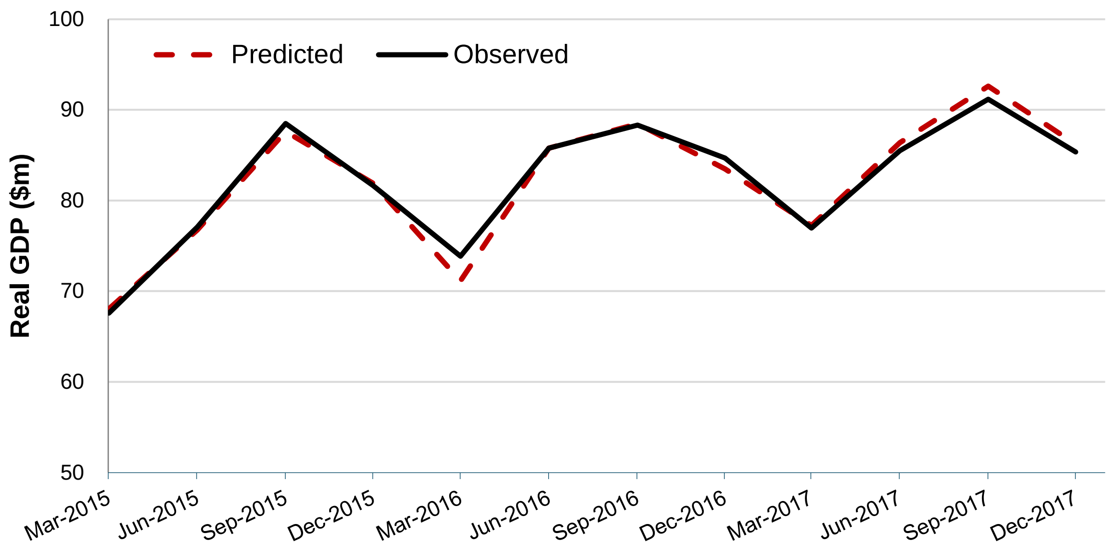

Source: MFEM analysis.

Table 9-2: Measures of static forecast accuracy, quarterly real GDP

Measures	Multivariate model
Mean error	-0.06
Root mean squared error	1.06
Mean absolute percentage error	0.97
Theil's U1	0.01

Source: MFEM analysis

10 Model performance against principles

The performance of the preferred model specification against the set of principles set out earlier in this attachment is summarised in Table 10-1.

Table 10-1: Performance against principles

Principle	Performance	Comments
Simple, transparent and replicable	✓	The forecasting approach comprises a standard multivariate ARIMA time series GDP model. The model development and specification has been described at length in this paper. A copy of the R script that contains the instructions for running the model and the model data can be made available on request.
Sound statistical basis	✓	The model utilises a common time series econometric method and the preferred specification has been subjected to a robust and objective statistical model selection process.
Forecast period	✓	The model has been configured to produce annual GDP forecasts from 2018–9 to 2021–22, covering the 4-year period forward budget period.
Forecast accuracy	✓	The model performed adequately when subjected to a formal out-of-sample forecast testing process. A combination forecast method is then applied using a simple average of 5 models.

11 Model forecasts

11.1 Introduction

The seventh and final step in the modelling process is to forecast real GDP, with quarterly forecasts summed by relevant financial year. The model is designed to produce an estimate of annual real GDP for 2018–19 and forecasts for each of remaining three years of the forward budget period. The estimate for 2018–19 is a combination of observed and modelled data.

11.2 Forecast combinations

Bates and Granger (1969) and Clemen (1989) found that combining multiple forecasts from several different methods on the same time series leads to increased forecast accuracy. Hyndman (2012) notes that:

> While there has been considerable research on using weighted averages, or some other more complicated combination approach, using a simple average has proven hard to beat.[23]

A combination approach is applied in this paper, utilising a simple average of five quarterly time series models – the multivariate seasonal ARIMA model described at length in this paper, and four univariate models:

- **Seasonal ARIMA model** – the 'auto.arima' function of the R 'forecast' package was applied to the real GDP time series to generate quarterly forecasts using an ARIMA (0,1,2)(1,1,1)[4] model.
- **Exponential smoothing or error, trend, seasonal (ETS) model** – the 'ets' function of the R 'forecast' package was applied to the real GDP time series to generate quarterly forecasts using an ETS (A,N,A) model.[24]

  Forecasts produced using exponential smoothing methods are weighted averages of past observations, with the weights decaying exponentially as the observations get older – i.e. the more recent the observation the higher the associated weight. Hyndman and Athanasopoulos (2012) state that this 'framework generates reliable forecasts quickly and for a wide spectrum of time series which is a great advantage and of major importance to applications in industry.'[25]
- **Seasonal naïve ARIMA model** – the 'snaive' function of the R 'forecast' package was applied to the real GDP time series to generate quarterly forecasts using an ARIMA (0,0,0)(0,1,0)[4] model. The seasonal naïve forecast is set equal to the last observed value from the same season of the year.

---

[23] Hyndman and Athanasopoulos, 2012: 12.4.

[24]

[25] Hyndman and Athanasopoulos, 2012: 7.

- **TBATS model** – an exponential smoothing state space model with a Box-Cox statistical transformation using ARMA (autoregressive and moving average) errors and incorporating a trigonometric representation of seasonal components.

11.3 Conditioning the forecast explanatory variables

11.3.1 Introduction

In order to produce GDP forecasts using the multivariate ARIMA model, assumptions are required about expected conditions (the levels of the explanatory variables) over the forecast period. The forecasts presented in this chapter are conditioned on the following assumptions about the explanatory variables.

11.3.2 GDP

A Holt-Winters seasonal exponential smoothing model was applied to the observed real GDP time series to generate forecasts for the GDP lag explanatory variable. The Holt-Winters seasonal method comprises the forecast equation and three smoothing equations — one for the level, one for the trend and one for the seasonal component.[26] The forecasts are shown in Figure 11-1.

Figure 11-1: Holt-Winters univariate GDP forecasts

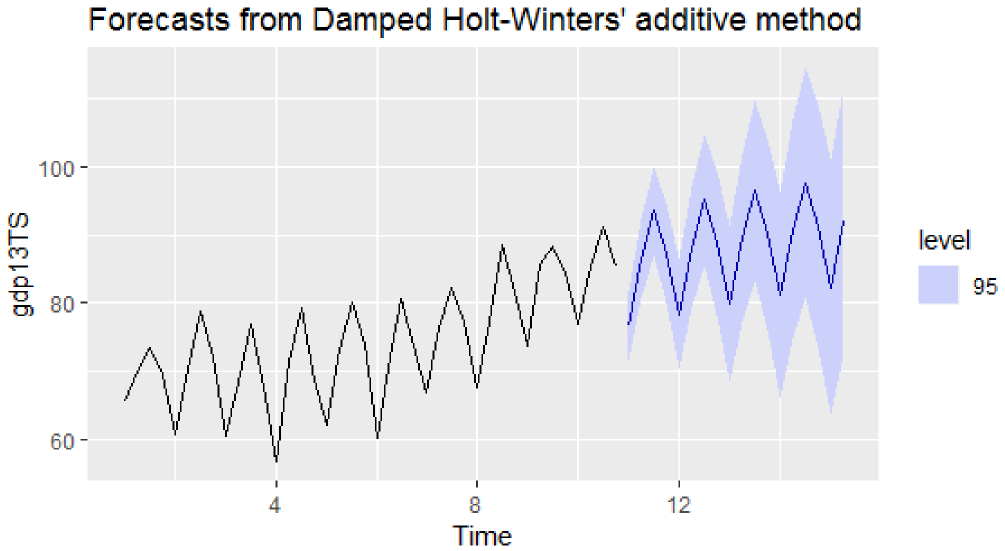

Source: MFEM analysis.

11.3.3 Consumer Price Index

The Holt-Winters approach was also applied to the observed CPI time series to generate forecasts for the CPI explanatory lags (see Figure 11-2).

[26] See Hyndman and Athanasopoulos (2012) section 7.3 for more detail on the Holt-Winters seasonal method.

Figure 11-2: Holt-Winters CPI forecasts

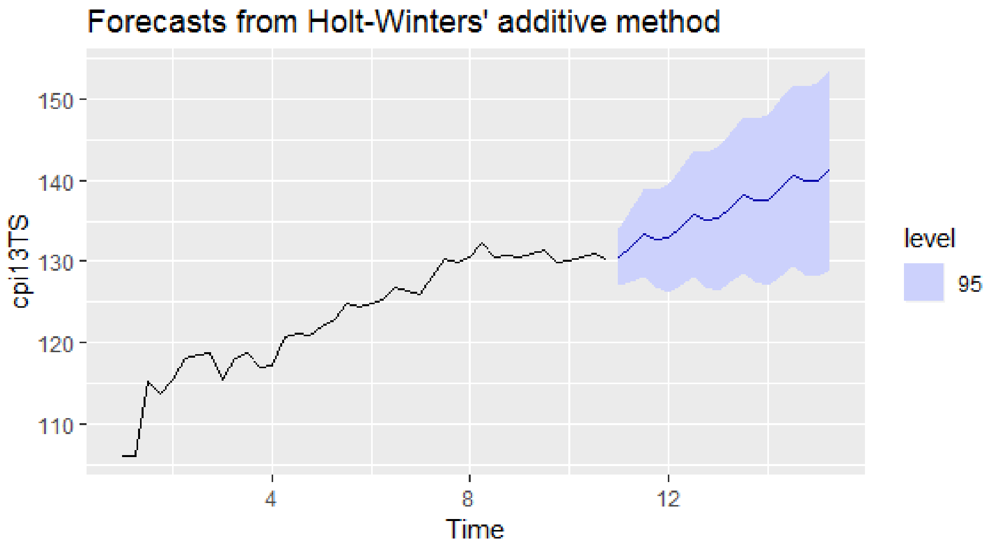

Source: MFEM analysis.

11.3.4 Imports

The Holt-Winters seasonal method was again applied to the observed imports time series to generate forecasts for the imports explanatory lags. As the seasonal variations are roughly constant throughout the series, the additive structure of the model was selected. In addition, the forecast was damped – using a damping coefficient – to more conservatively estimate the predicted trend. This method results in constant long-term forecast without affecting the trend of short-term forecasts.

Figure 11-3: Holt-Winters imports forecasts

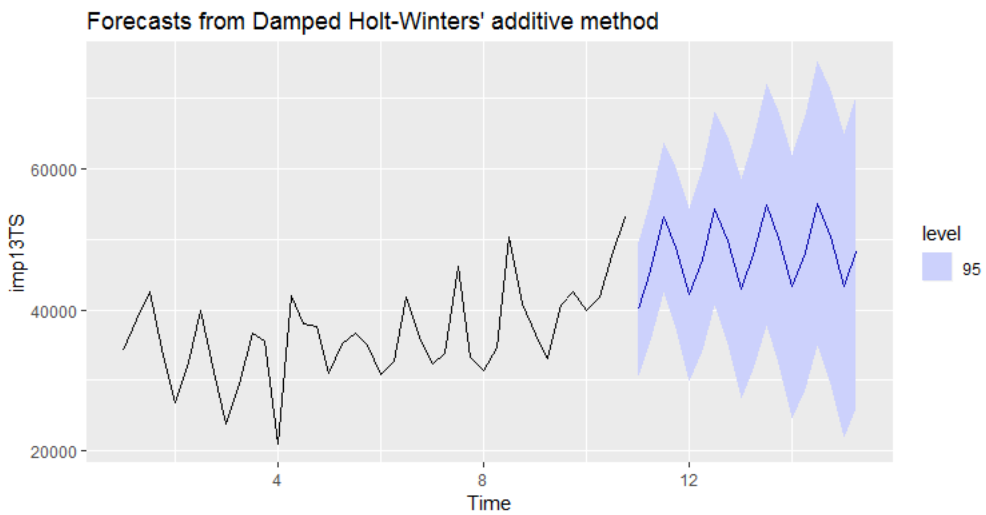

Source: MFEM analysis.

11.3.5 Visitor arrivals

As discussed in Chapter 5, the Cook Islands Statistics Office publishes monthly arrivals data with a time series available from January 1995 to June 2018. This monthly time series has been used to forecast arrivals for the purposes of the arrivals lags required for the multivariate model. The Holt-Winters seasonal method described above for the imports forecasts was applied to generate monthly arrivals forecasts, subsequently aggregated into quarterly data (see Figure 11-4).

Figure 11-4: Holt-Winters arrivals forecasts, monthly

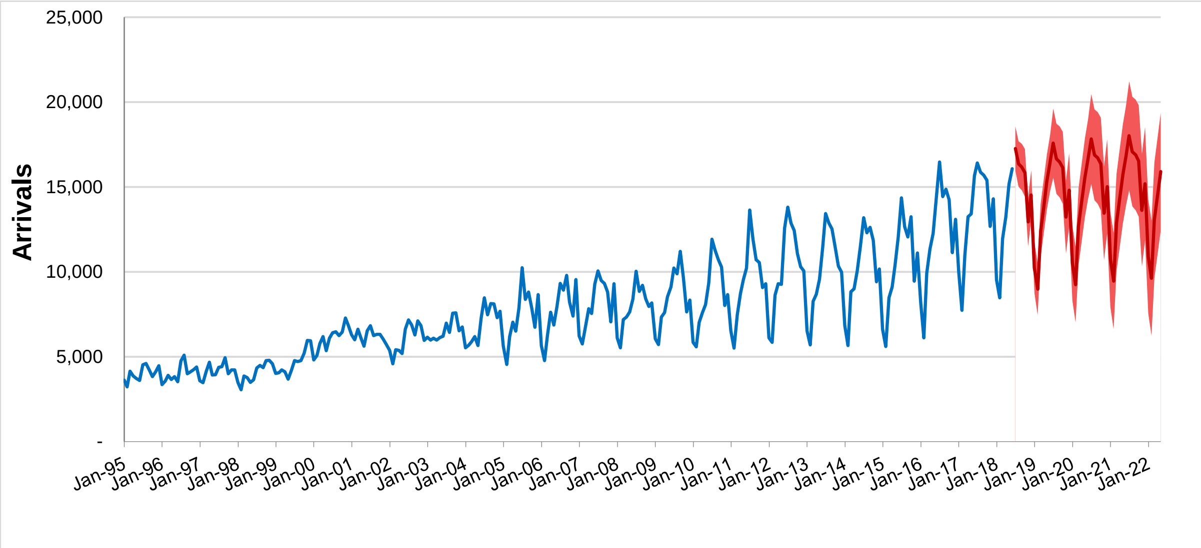

Source: MFEM analysis.

In line with the damped approach adopted, the arrivals forecasts show a declining trend over the forecast period, as shown in Figure 11-5.

Figure 11-5: Arrivals forecasts, annual percentage change

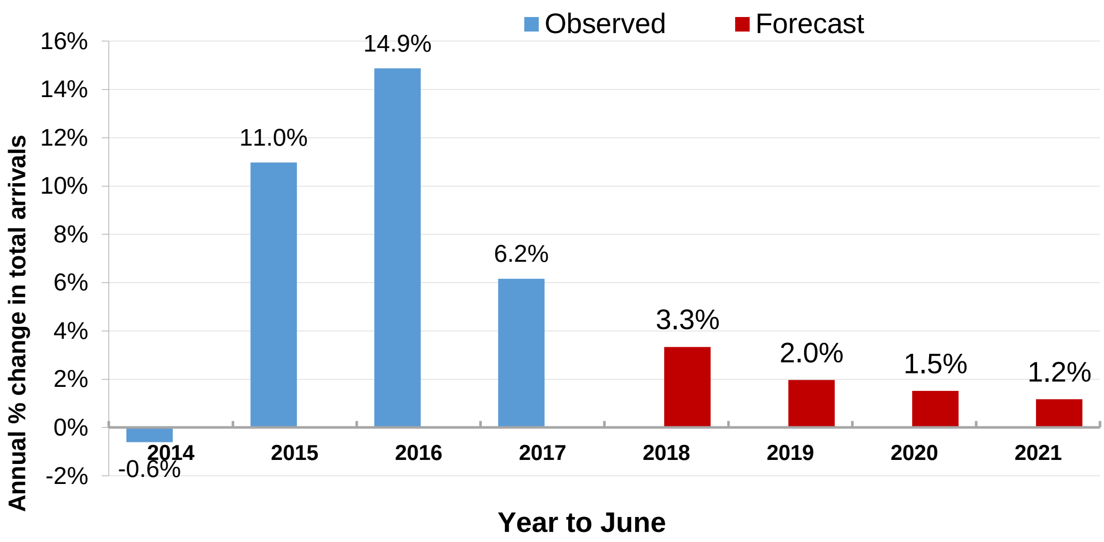

Source: MFEM analysis.

11.4 GDP forecasts

Observed and forecast quarterly real GDP from December 2017 to June 2022 is shown in Figure 11-6. A 95 per cent high and low confidence interval has also been computed and is displayed as the shaded area either side of the point forecast time series. Commenting on prediction intervals for multivariate models, Hyndman and Athanasopoulos (2012) make a salient point:

It is important to realise that the prediction intervals from regression models (with or without ARIMA errors) do not take into account the uncertainty in the forecasts of the predictors. So they should be interpreted as being conditional on the assumed (or estimated) future values of the predictor variables.[27]

Figure 11-6: Observed and forecast real GDP, by quarter

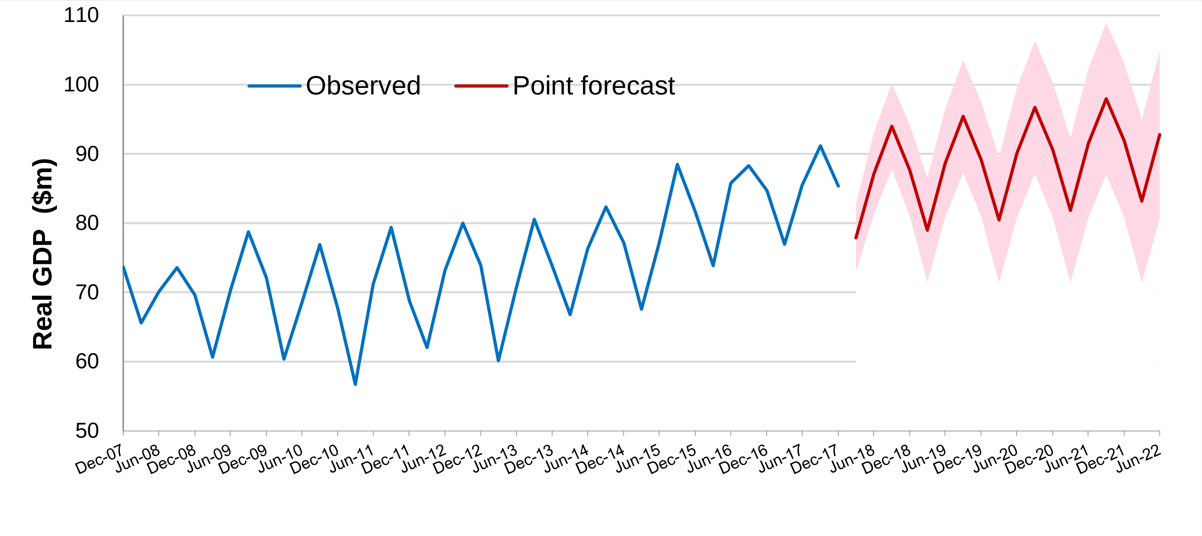

Source: MFEM analysis.

Table 11-1 shows the quarterly forecasts aggregated into the financial year forecasts required for the forward budgeting period. The forecast percentage change over the forward period compared to recent years is shown in Figure 11-7. A growth rate of 1.8 per cent is estimated in 2017–18, with growth forecast to rise to 2.4 per cent in 2018–19 and then fall to 1.4 per cent by the final year of the forward budget period. Average annual growth over the 4-year forward period is forecast at 1.7 per cent.

[27] Hyndman and Athanasopoulos, 2012: 9.3.

Table 11-1: GDP forecasts to 2022, $m

Year to	GDP real ($m)	% change	Implicit GDP price deflator	GDP nominal ($m)	% change
Jun-14	297.4		127.77	380.0	
Jun-15	304.1	2.3%	129.18	392.9	3.4%
Jun-16	329.8	8.4%	125.87	415.1	5.6%
Jun-17	335.5	1.7%	126.18	423.3	2.0%
Jun-18e	341.5	1.8%	127.16	434.3	2.6%
Jun-19f	349.3	2.3%	127.62	445.7	2.6%
Jun-20f	355.2	1.7%	128.31	455.7	2.2%
Jun-21f	360.6	1.5%	128.83	464.6	1.9%
Jun-22f	365.8	1.4%	129.45	473.5	1.9%

Source: MFEM analysis.

Figure 11-7: Annual change in real GDP, percentage

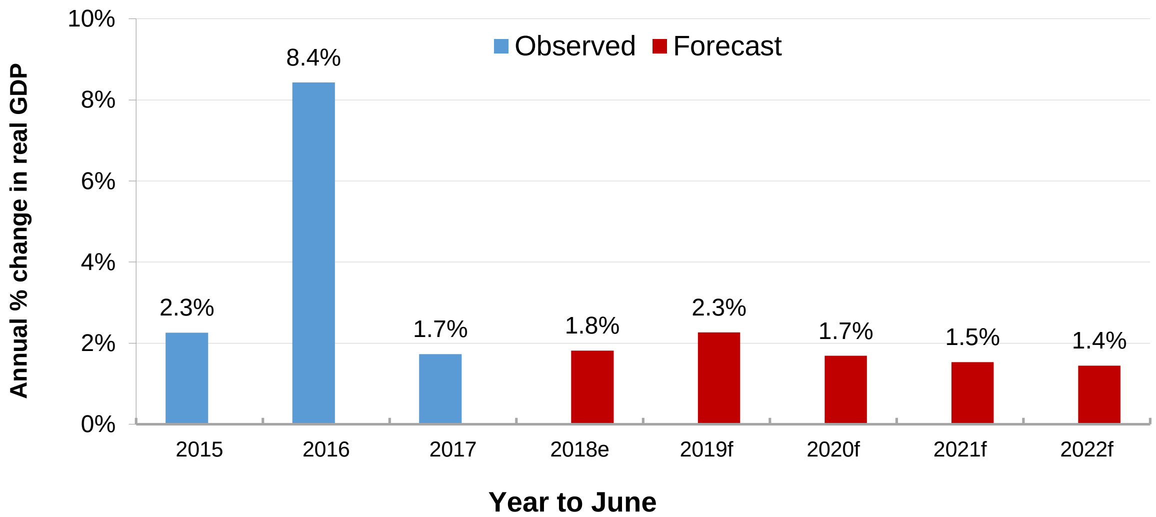

Source: MFEM analysis.

Abbreviations and acronyms

ABS	Australian Bureau of Statistics
ACF	Autocorrelation function
AIC	Akaike Information Criterion
AR	Autoregressive
ARIMA	Autoregressive integrated moving average
GDP	Gross Domestic Product
HES	Household Expenditure Survey
MFEM	Ministry of Finance and Economic Management
OLS	Ordinary least squares
PACF	Partial autocorrelation function

References

Bates, J. M., & Granger, C. W. J. (1969). The combination of forecasts. *Operational Research Quarterly*, *20*(4), 451–468. https://doi.org/10.1057/jors.1969.103

Box, George E. P., and G. M. Jenkins. 1970. *Time Series Analysis: Forecasting and Control*. Wiley.

Clemen, R. (1989). Combining forecasts: A review and annotated bibliography with discussion. *International Journal of Forecasting*, *5*, 559–608.

Cook Islands Government. 2017. "Budget Estimates 2017/18 Book 1: Appropriation Bill, Appropriations and Commentary." Rarotonga.

Elliot, Graham, and Allan Timmerman. 2008. "Economic Forecasting." *Journal of Economic Literature* 46 (1): 3–56.

Granger, CWJ 1986. "Developments in the study of cointegrated variable." Oxford Bulletin of Economics and Statistics. 48, 3, 213-228.

Greene, William H. 2012. *Econometric Analysis*. 7th edition. New York: Prentice Hall.

Hyndman, R. J. (2001). Box-Jenkins modelling.

Hyndman, RJ and Athanasopoulos, G 2012. "Forecasting: principles and practice: Online textbook on forecasting." https://www.otexts.org/fpp.

Hyndman, RJ and Yeasmin Khandakar. 2008. "Automatic Time Series Forecasting: The Forecast Package for R." *Journal of Statistical Software* 27: 1–22.

Kwiatowski, D., P. C. B. Phillips, et al. (1992). "Testing the null hypothesis of stationarity against the alternative of a unit root." Journal of Econometrics **54**: 159-178.

Pfaff, B. (2008). Analysis of Integrated and Cointegrated Time Series with R. New York, Springer.

Phillips, P. C. B. and S. Ouliaris (1990). "Asymptotic Properties of Residual Based Tests for Cointegration." Econometrica **58**(1): 165-193.

R Core Team (2016). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria.

RStudio Team (2015). RStudio: Integrated Development for R. RStudio, Inc. Boston, MA.