Group 2 ECON2006 COURSEWORK

Word count: **1988** (excluding graphs, tables and appendices)

NAME	**STUDENT ID**
George Bennett	[redacted]
Benjamin Hunter	[redacted]
Oliver Chapman	[redacted]
Michael Nixon	[redacted]
George Williams	[redacted]

Contents

Introduction…… 3

Question 1…… 3 - 5

Question 2…….. 6 – 7

Question 3…… 8 - 12

Question 4…… 13

Appendix…….. 14 - 17

Introduction:

During this investigation, we utilise econometric time series analysis through Stata to analyse whether the inclusion of lagged values for inflation enhances forecasting for the change in unemployment in Japan and Canada.

Task 1:

Summary Statistics:

	Obs	Mean	Std. Dev.	Min	Max
dinf CAN	279	0.005	0.838	-4.897	3.307
dunemp CAN	259	0.003	0.355	-1.000	1.733
dinf JPN	255	0.005	1.166	-8.236	7.233
dunemp JPN	259	-0.001	0.120	-0.300	0.533

Figure 1

Japan and Canada both have relatively low and stable changes in the mean year-on-year inflation change. A standard deviation of 1.1657 suggests higher volatility compared to Canada, with a standard deviation of 0.8381. The minimum and maximum values were -8.24 and 7.23 respectively, which are larger than the values for Canada (-4.90 and 3.31). Changes in inflation in Japan were more extreme than in Canada, possibly because Japan was more sensitive to economic shocks, or inflation was better contained in Canada.

Changes in unemployment in Japan are highly stable, with a mean change close to zero (−0.00077) and a much lower standard deviation (0.12 vs 0.36 in Canada). Canada's wider range of unemployment changes suggests a more flexible labour market that responds more strongly to economic cycles.
Overall, the summary statistics suggest that both variables display considerable variation over time, which provides scope for econometric modelling of their relationship.

Time Series Properties:

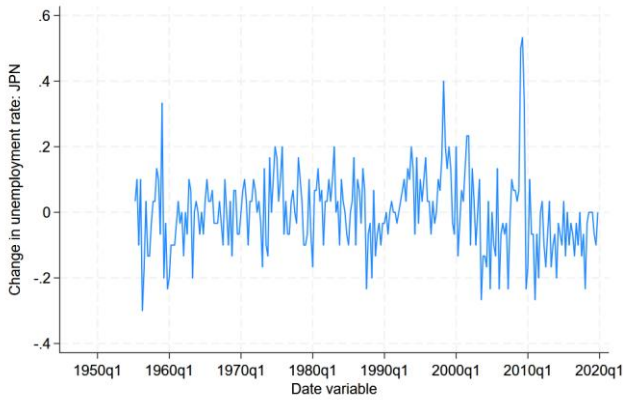


Figure 2

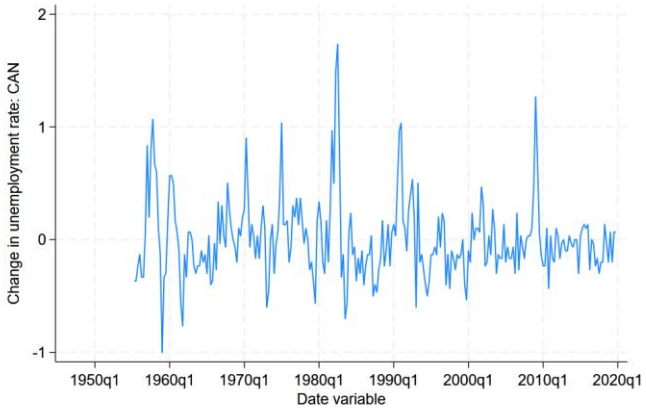


Figure 3

The change in unemployment series for Canada (*Figure 2*) fluctuates around a mean without a clear deterministic trend, suggesting stationarity. Nevertheless, noticeable spikes correspond to major economic downturns. For example, sharp increases in unemployment changes are visible during recessions, reflecting the cyclical nature of labour market conditions. Likewise, the changes in unemployment in Japan (*Figure 3)*, exhibit stationarity, and changes fluctuate between positive and negative values.

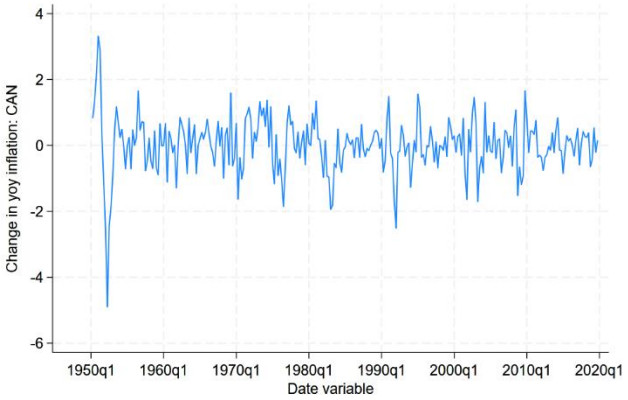


Figure 4

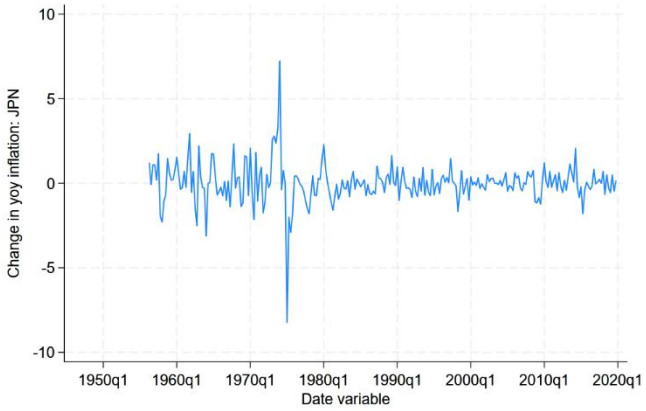


Figure 5

Figure 4 and *Figure 5* also exhibit stationarity, with changes in inflation fluctuating between positive and negative values around a constant mean.
These graphs suggest that both variables exhibit short-term fluctuations rather than long-run trends, supporting the use of autoregressive models to capture their dynamics.

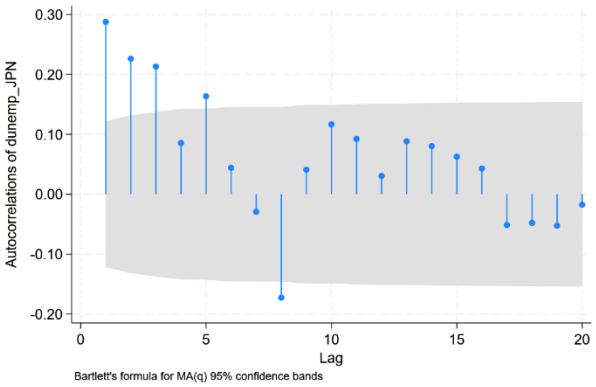


Figure 6

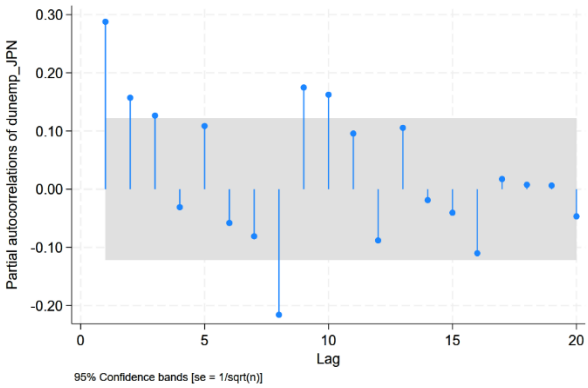


Figure 7

The ACF for Japan (*Figure* 6) shows statistically significant positive autocorrelation at lags 1,2,3, and 5, suggesting the labour market in Japan exhibits some persistence, potentially indicating that a change in unemployment in one period is positively correlated with changes in the following periods. Negative correlation at lag 8 suggests a potential cyclical effect.

The PACF for Japan (*Figure 7)* shows statistically significant partial autocorrelations at lags 1,2,3,9 and 10, as well as a negative spike at lag 8. This supports the findings of the ACF correlogram. A drop off after lag 3 suggests an AR(3) model could be most appropriate for this data.

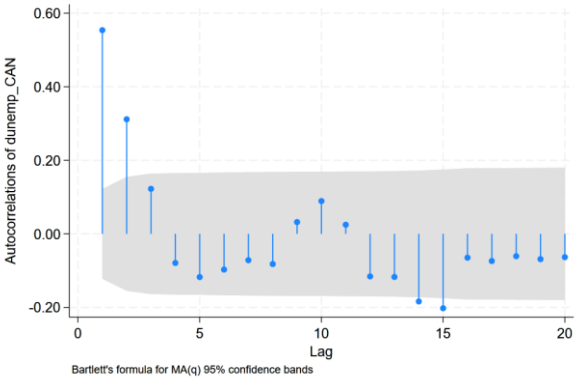


Figure 8

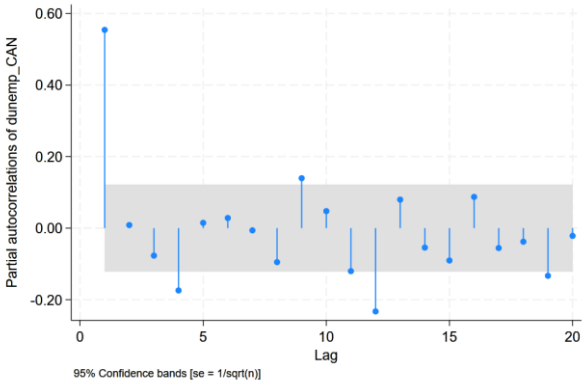


Figure 9

The ACF for Canada (*Figure 8)* indicates that the correlation between current and past values declines relatively quickly as the lag length increases. The PACF (*Figure* 9) shows that the first lag appears particularly important, while the significance of further lags diminishes. This pattern suggests that changes in unemployment are likely influenced by their own recent past values.

Significant lags in the ACF and PACF mean we should use a formal model selection process to determine the best-fitting model to analyse this relationship.

Task 2:

Our investigation generated and compared 20 ARDL and AR models, reporting their AIC values and p-values of the Breusch-Godfrey (BG) tests through our in-sample period from 1955q2 – 2015q4. Comparison table located in the appendix.

Methodology:

This investigation has selected its favoured models via a two-step process:

1. Diagnostic Filtering through Breusch-Godfrey Testing

This stage eliminates models that contain serial correlation; in turn, removing models that contain errors correlated over time and avoiding inaccurate estimation of standard errors, which could lead to invalid inferences being made due to skewed t-statistics. This could compromise forecasts, damaging the integrity of the report.

To screen serial correlation out of our investigation, we conducted a BG test at the 5% level of significance at lag order 4, which allows the investigation to observe for seasonality over the year. This test was chosen as it includes lagged dependent variables and allows for a higher-order serial correlation. In this test, the null hypothesis assumes that there is no serial correlation. As a result, we rejected the null and screened out models with p-values smaller than 0.05. These models were deemed by the test to contain serial correlation.

2. Identifying the primary candidate models

Based on our remaining models that have not been screened out from the dataset, the investigation has chosen the two ARDL models which have the lowest AIC; these have the maximum amount of information and the minimum number of variables. We have also included models for comparison for both Japan and Canada.

Selected models for Japan:

Model	p	q	AIC	bg_p	N	keep
ARDL	3	1	-347.387	0.932062	238	1
ARDL	4	1	-345.751	0.946203	238	1
AR	3	0	-344.066	0.033021	240	0

Figure 10

Figure 10 displays the 3 chosen models for Japan in this investigation. The favoured models for Japan are ARDL (3,1), and ARDL (4,1), with AR (3) for comparison.

Both these models passed the initial screening process as they demonstrate no evidence of serial correlation, with p-values of 0.93 and 0.95, respectively, from the BG test. These models also hold the lowest AIC scores, which indicates the best in-sample fit for our investigation.

We have then chosen a model for comparison - AR (3). This model is strong for comparison as it ignores the effect of inflation (q = 0), helping to investigate the effect on performance caused by the inclusion of lagged inflation.

AR (3) failed in the serial correlation screening stage of the investigation because its p-value (0.033) is smaller than 0.05. Indicating that there is statistically significant evidence of serial correlation in this model, making its inferences invalid.

Comparing the chosen AR model to our favoured models, such as ARDL (3,1), we can infer that excluding inflation may have induced serial correlation, suggesting omitted-variable bias through correlated errors. This indicates the relevance of inflation to unemployment in the Japanese economy.

Models for Canada:

Model	p	q	AIC	bg_p	N	keep
ARDL	1	4	114.1233	0.251196	242	1
ARDL	1	2	114.1258	0.205951	242	1
ARDL	3	1	120.8365	0.020894	240	0

Figure 11

Figure 11 shows our preferred models for Canada ARDL (1,4) and ARDL (1,2), with ARDL (3,1) being the model for comparison.

The ARDL (1,4) model passes the BG test for serial correlation, indicating that the residuals are uncorrelated, suggesting that the model sufficiently captures the dynamic behaviour of unemployment changes.

Our second model is the ARDL (1,2) model, which also performs well, displaying a slightly higher AIC value but similarly passing the serial correlation test. This model, therefore, represents an alternative specification with strong empirical performance.

For comparison purposes, a third model was selected from the remaining specifications. The ARDL (3,1) model was chosen as it has a relatively high AIC value and therefore represents a weaker specification relative to the preferred models.

The results suggest that models incorporating lagged inflation changes provide a better description of unemployment dynamics in Canada. Since the ARDL models have lower AIC scores than autoregressive models, we can infer that inflation changes may contain useful predictive information for unemployment changes.

Task 3:

Point estimates: Japan

Using the three models selected previously, we generated out-of-sample point forecasts and 90% interval forecasts. Each model was estimated on the in-sample period up to and including 2015q4, and the estimated coefficients were used to produce dynamic forecasts of the change in unemployment for Canada and Japan. Quarterly dummies were included in all specifications to account for seasonal variation.

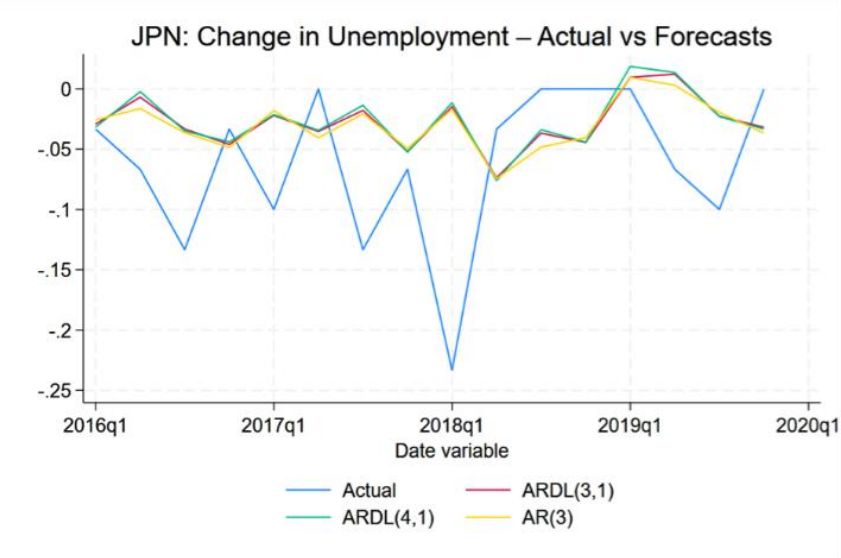


Figure 12

Figure 12 illustrates the 3 selected models' one-step-ahead point forecasts versus the actual figures for unemployment change over the out-of-sample period (2016q1 – 2019q4) for Japan.

All selected models demonstrate similar movements as they contain similar lagged variables. These tend to track the overall direction of actual data well but systematically underpredict movements in the actual data. The movements are conservative, suggesting mean reversion. Possibly hinting that there are variables that the models have not captured.

Point forecasts: Canada

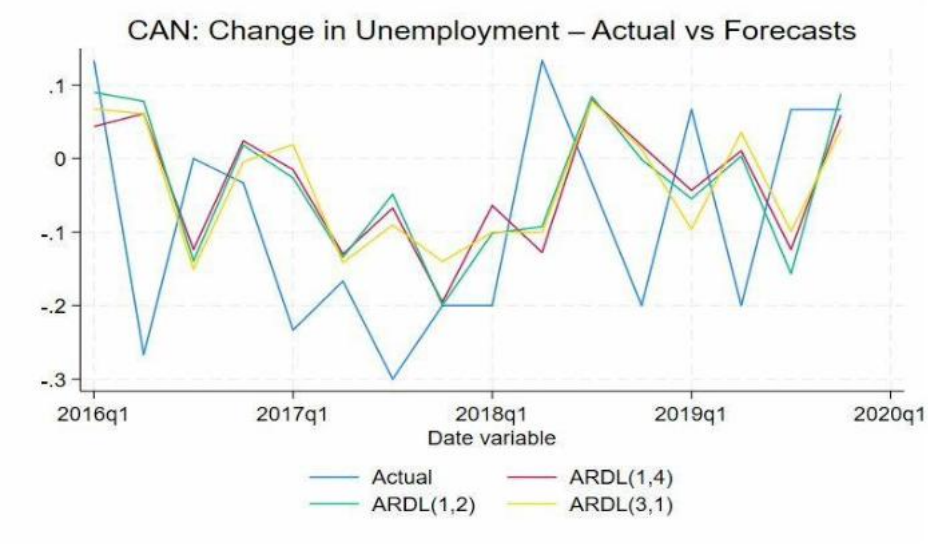


Figure 13

Figure 13 plots the realised changes in unemployment alongside the one-step-ahead point forecasts from all three models over the out-of-sample period (2016q1 – 2019q4).

First, all three models generate forecasts that are broadly similar to one another. This is expected given that all specifications share at least one lag of unemployment.

Additionally, the actual series is considerably more volatile than the forecasts. The models systematically underpredict the size of swings in the actual data, producing mean-reverting forecasts that fail to track the large quarter-to-quarter movements visible in the data.

Overall, the models for Japan and Canada capture average dynamics well but fail to anticipate sharp, idiosyncratic shocks.

Forecast intervals: Japan

In the next part of this investigation, we construct 90% forecast intervals around the point forecasts of each selected model, indicating the range within which future values have a 90% probability of falling.

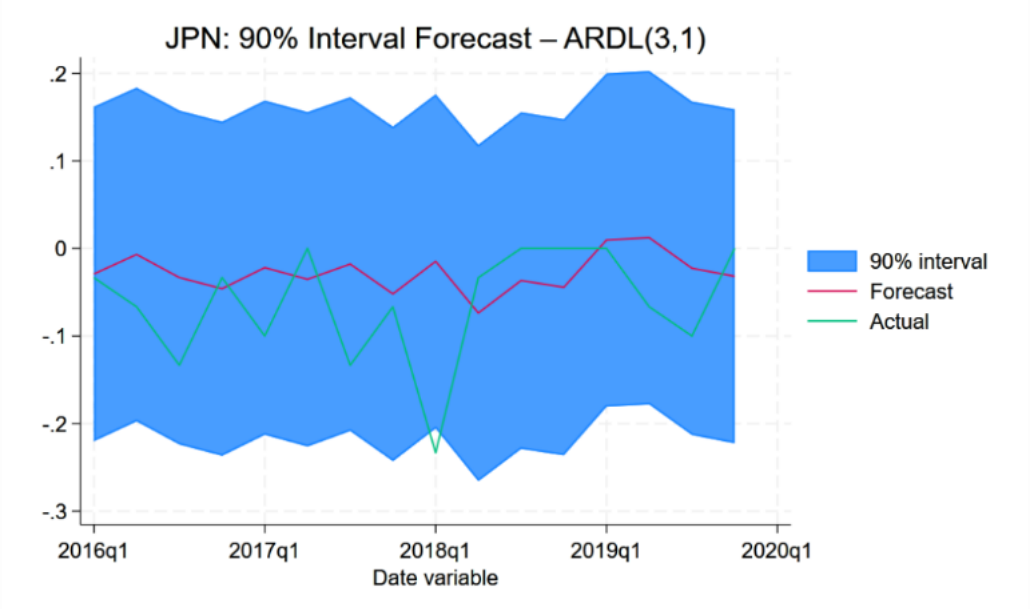


Figure 14

Figure 14 shows the preferred model from the previous section is ARDL (3,1). It has a broad interval of around ±0.2 percentage points around the point forecast, indicating that there is a high degree of forecast uncertainty. The actual values only exceeded the forecast interval in 2018q1, which indicates that the forecast is reliable at the 90% level.

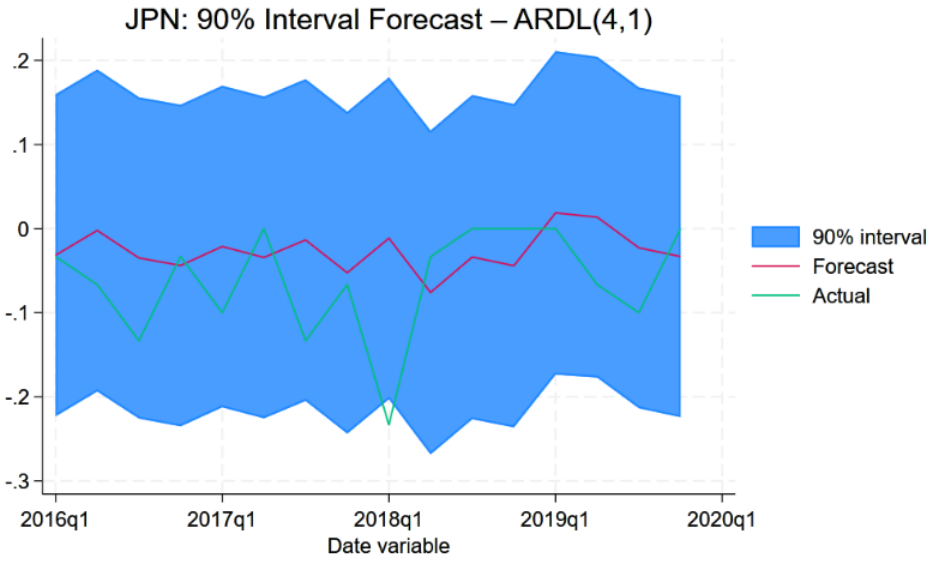


Figure 15

Here, *Figure 15* illustrates our second model of preference (ARDL (4,1)). This forecast has a similar breadth to the forecast interval for ARDL (3,1), which indicates a high degree of forecast uncertainty. Again, the actual value only exceeds the forecast interval in 2018q1, indicating that this forecast is reliable at the 90% level.

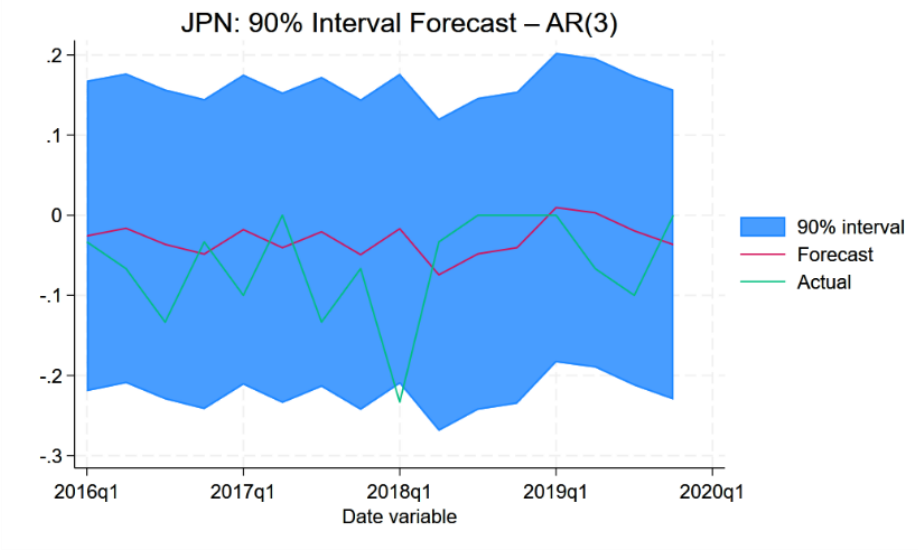


Figure 16

Model AR (3,0) is the model the investigation has used for comparison. *Figure 16* shows this model demonstrates a very similar interval to the previous two models. This indicates a high degree of forecast uncertainty. The actual values only exceed the interval in 2018q1, indicating that the forecast is reliable at the 90% level.

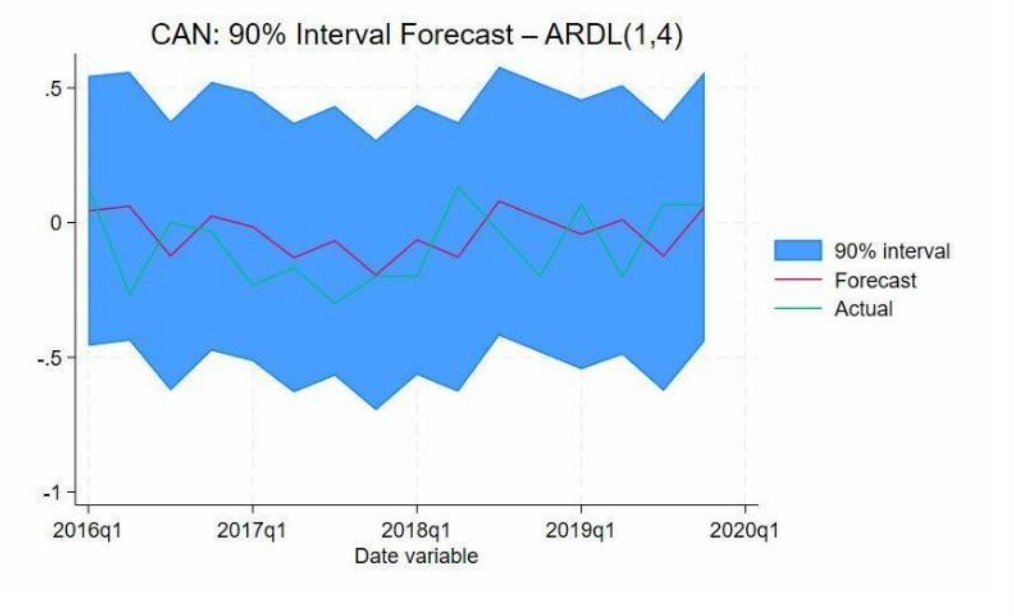


Figure 17

Figure 17 presents the 90% forecast interval for the preferred model, ARDL (1,4), together with the point forecast and the actual series over the out-of-sample period. The interval is broad, spanning ±0.5 percentage points around the point forecast, illustrating a high level of forecast uncertainty. Despite this width, the interval is reliable and captures all variation in the actual series.

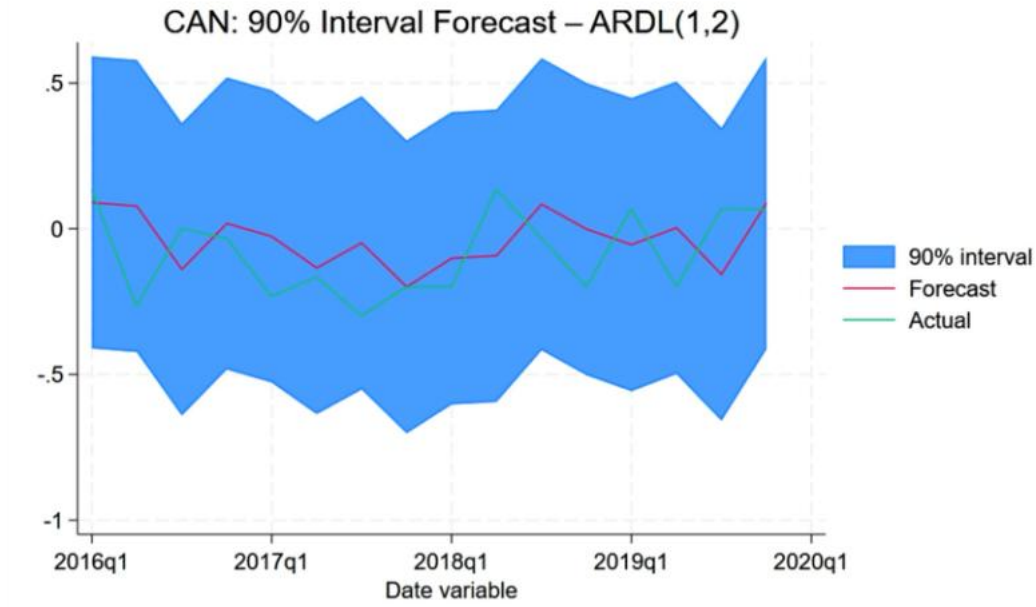


Figure 18

Figure 18 presents the 90% forecast interval for the change in Canadian unemployment over the out-of-sample period. Notably, all observed outcomes fall within the 90% forecast interval, suggesting the model generates reliable probabilistic forecasts.

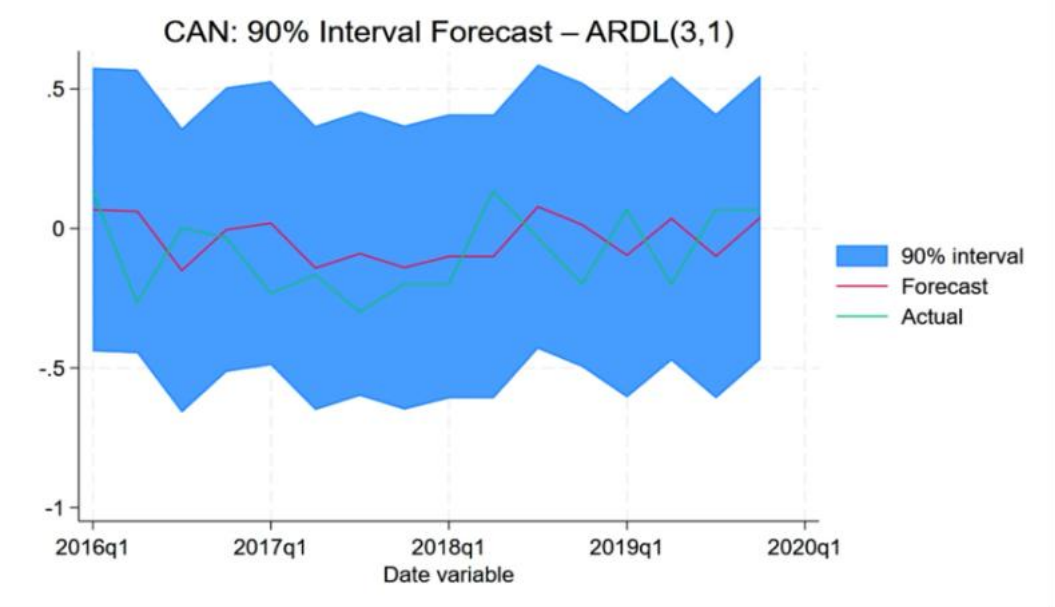


Figure 19

Figure 19 shows all observed values remain within the 90% forecast interval, indicating that the model adequately characterises forecast uncertainty. Relative to the ARDL (1,2), the ARDL (3,1) produces marginally more stable point forecasts; the two specifications deliver broadly comparable out-of-sample performance overall.

TASK 4:

To evaluate the three out-of-sample models by their predictive power, we use the Root Mean Squared Error (RMSE)

$$RMSE = \sqrt{\frac{1}{m}\sum_{h=1}^{m} e_{n+h}^2}$$

In our model comparison, we chose the model with the smallest RMSE value to maximise predictive power. The smallest RMSE value represents the smallest difference between the predicted and actual unemployment values.

The results for the chosen models are shown below:

Canada:

Model	***RMSE Value***	***Ranking***	***AIC Value***	***Ranking***
ARDL(1,2)	0.1714	1st	114.1258	2nd
ARDL(1,4)	0.1726	2nd	114.1233	1st
ARDL(3,1)	0.1731	3rd	120.8365	3rd

Figure 20

Japan:

Model	***RMSE Value***	***Ranking***	***AIC Value***	***Ranking***
ARDL(3,1)	0.07934	2nd	-347.387	1st
ARDL(4,1)	0.08057	3rd	-345.751	2nd
AR(3,0)	0.07863	1st	-344.066	3rd

Figure 21

Figure 20 and *Figure* 21 show that the Canada ARDL (1,2) and Japan AR (3) models have the lowest RMSE values (0.1714 and 0.07863 respectively).

For Canada, two of the models have very similar performance. ARDL(1,2) has the lowest RMSE, but a slightly higher AIC than ARDL(1,4). Both models perform well. We must decide whether we value RMSE or AIC more in our chosen model. We recommend ARDL (1,2), as it has the lowest RMSE value and is therefore more accurate in forecasting unemployment, despite the lost efficiency via a higher AIC. We have decided the accuracy of forecasts is more important.

For Japan, the AR(3) model produces the lowest RMSE of 0.07863. However, this model failed the Breusch-Godfrey test ($p = 0.033$), indicating serial correlation in the residuals, which renders its inferences unreliable. The absence of inflation lags likely induced this misspecification through omitted variable bias. The ARDL(3,1) model is therefore recommended: it has the second-lowest RMSE while passing the BG test and capturing the distributed lag effect of inflation on unemployment.

In summary, we recommend ARDL(1,2) for Canada and ARDL(3,1) for Japan — both pass diagnostic testing, minimise forecast error, and capture the inflation-unemployment relationship.

Appendix:

Regression Output for Question 2 – Japan:

	(1) Change in unemployment rate: JPN	(2) Change in unemployment rate: JPN	(3) Change in unemployment rate: JPN
L.Change in unemployment rate: JPN	0.272***	0.271***	0.224***
	(0.0642)	(0.0642)	(0.0649)
L2.Change in unemployment rate: JPN			0.118
			(0.0660)
L3.Change in unemployment rate: JPN			0.129*
			(0.0650)
L.Change in yoy inflation: JPN	-0.0111	-0.00976	
	(0.00670)	(0.00664)	
L2.Change in yoy inflation: JPN	-0.00245	-0.000207	
	(0.00676)	(0.00658)	
L3.Change in yoy inflation: JPN	0.000848		
	(0.00675)		
L4.Change in yoy inflation: JPN	0.0134*		
	(0.00662)		
Constant	0.0158	0.0132	0.0139
	(0.0152)	(0.0152)	(0.0151)
Observations	235	237	240
AIC	-336.9	-339.2	-344.1
BG p-value	0.0864	0.0598	0.0330

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Regression Output for Question 2 – Canada:

	(1) Change in unemployment rate: CAN	(2) Change in unemployment rate: CAN	(3) Change in unemployment rate: CAN
L.Change in unemployment rate: CAN	0.558***	0.563***	0.554***
	(0.0539)	(0.0538)	(0.0654)
L2.Change in unemployment rate: CAN			0.0458
			(0.0748)
L3.Change in unemployment rate: CAN			-0.0872
			(0.0668)
L.Change in yoy inflation: CAN	-0.0247	-0.0236	-0.0168
	(0.0283)	(0.0283)	(0.0289)
L2.Change in yoy inflation: CAN	0.0650*	0.0666*	
	(0.0289)	(0.0282)	
L3.Change in yoy inflation: CAN	-0.0157		
	(0.0288)		
L4.Change in yoy inflation: CAN	0.0554		
	(0.0282)		
Constant	0.0166	0.0175	0.0170
	(0.0388)	(0.0390)	(0.0395)
Observations	242	242	240
AIC	114.1	114.1	120.8
BG p-value	0.251	0.206	0.0209

Standard errors in parentheses
* p < 0.05, ** p < 0.01, *** p < 0.001

ARDL Models for Question 2 - Japan:

Model	p	q	AIC	bg_p	N	Keep
ARDL	3	1	-347.3868729	0.932061939	238	1
ARDL	4	1	-345.7514148	0.946202831	238	1
ARDL	2	1	-344.8155248	0.320499605	238	1
ARDL	3	2	-343.7228574	0.961192147	237	1
AR	2	0	-343.3653831	0.154590932	241	1
ARDL	1	1	-343.3186675	0.062520412	238	1
ARDL	3	4	-342.4847291	0.59835442	235	1
ARDL	4	2	-341.9730032	0.894986388	237	1
ARDL	2	2	-341.0189702	0.415446732	237	1
ARDL	3	3	-340.8894492	0.377208435	236	1
ARDL	4	4	-340.506925	0.576571643	235	1
ARDL	2	4	-339.2859586	0.264830885	235	1
ARDL	1	2	-339.1942593	0.059805607	237	1
ARDL	4	3	-338.9868925	0.179688778	236	1
ARDL	1	4	-336.9386469	0.086442069	235	1
ARDL	1	3	-335.5019124	0.056164613	236	1
AR	3	0	-344.0662074	0.03302111	240	0
AR	1	0	-341.7430792	0.026575858	242	0
AR	4	0	-340.7437701	0.012721833	239	0
ARDL	2	3	-338.0886573	0.038415846	236	0

ARDL Models for Question 2 - Canada:

Model	p	q	AIC	bg_p	N	Keep
ARDL	1	4	114.1233311	0.251196056	242	1
ARDL	1	2	114.1258445	0.205951203	242	1
AR	4	0	114.4293496	0.59387008	239	1
ARDL	4	2	114.642766	0.772614017	239	1
ARDL	4	4	115.4969746	0.846729297	239	1
ARDL	4	1	115.8413301	0.976276778	239	1
AR	1	0	115.8768838	0.107224462	242	1
ARDL	1	3	116.1067688	0.207751719	242	1
ARDL	2	4	116.4027471	0.236233006	241	1
ARDL	2	2	116.4255607	0.19740849	241	1
ARDL	4	3	116.5179149	0.809082378	239	1
ARDL	3	2	117.5532807	0.059119349	240	1
ARDL	1	1	117.7805117	0.078947367	242	1
ARDL	3	4	118.0593831	0.220158383	240	1
AR	2	0	118.1528487	0.079178426	241	1
ARDL	2	3	118.402501	0.189553111	241	1

AR	3	0	119.1843653	0.095664399	240	1
ARDL	3	3	119.5049406	0.069808333	240	1
ARDL	2	1	120.0614124	0.067757237	241	1
ARDL	3	1	120.8364917	0.020894022	240	0

Table of Mean Forecast Errors – Japan:

Table of Mean Forecast Errors - JPN

	mean
Absolute value of Forecast Errors from the ARDL(3,1) model	0.0598725
Absolute value of Forecast Errors from the ARDL(4,1) model	0.0608641
Absolute value of Forecast Errors from the AR(3) model	0.0603469
Square of Forecast Errors from the ARDL(3,1) model	0.0062943
Square of Forecast Errors from the ARDL(4,1) model	0.0064909
Square of Forecast Errors from the AR(3) model	0.0061822
Observations	16

Table of Mean Forecast Errors – Canada:

Table of Mean Forecast Errors - CAN

	mean
Absolute value of Forecast Errors from the ARDL(1,4) model	0.1460618
Absolute value of Forecast Errors from the ARDL(1,2) model	0.1425076
Absolute value of Forecast Errors from the ARDL(3,1) model	0.1480652
Square of Forecast Errors from the ARDL(1,4) model	0.0298004
Square of Forecast Errors from the ARDL(1,2) model	0.0293906
Square of Forecast Errors from the ARDL(3,1) model	0.0299755
Observations	16