Forecasting economic downturns in South Africa using leading indicators and machine learning

Fourie, Jurgens and Steenkamp, Daan

Codera Analytics, Codera Analytics. Research Associate with the Economics Department at Stellenbosch University

7 May 2025

Forecasting economic downturns in South Africa using leading indicators and machine learning

Jurgens Fourie[*] Daan Steenkamp[†]

May 7, 2025

Abstract

We identify South African business cycles using the algorithm of Bry-Boschan and show that the identified turning points are very similar to those from other approaches. We demonstrate that South Africa has a very volatile business cycle that makes it particularly difficult to predict turning points in the economic cycle. South Africa's business cycle is characterised by relatively long downswings and short upswing phases with low amplitude. We find that the South African Reserve Bank (SARB)'s Leading Indicator does not substantive improve predictions of the business cycle relative to GDP itself. We assess the performance of a range of potential leading indicators in identifying economic downturns and consider whether alternative indicators and estimation approaches can produce better predictions than those of the SARB. We demonstrate that using a larger information set produces substantially better business cycle predictions, especially when using machine learning techniques. Our findings have implications for the creation of composite leading indicators, with our results suggesting that many of the macroeconomic variables considered by analysts as leading indicators do not provide good signals of GDP growth or developments in the South African business cycle.

JEL classification: E32, E37

Keywords: business cycle, forecast, leading indicator, economic downturns

[*]Codera Analytics.

[†]Codera Analytics. Research Associate with the Economics Department at Stellenbosch University. *Corresponding author.* Email: daan@codera.co.za

1 Introduction

Understanding business cycles is crucial for economic forecasting and monetary policy as early signals of where the economy is heading enables industry and policymakers to be more proactive in their strategic planning.

In this policy paper, we identify business cycles using the algorithm of (Bry and Boschan, 1971) as adapted by (Harding and Pagan, 2002). The Bry-Boschan (BBQ) algorithm detects peaks and troughs by applying minimum criteria for phase durations (trough-to-peak or peak-to-trough) and complete cycles (an upswing followed by a downswing).[1] Figure 1 illustrates the intuition of the approach for identifying business cycles using a Hodrick-Prescott filter and the BBQ algorithm.

We compare our business cycle estimates to other South African estimates from the South African Reserve Bank (SARB, see Venter and Pretorius (2004), Venter (2020), Venter and Wolhuter (2023)) and Organisation for Economic Co-operation and Development (OECD). Although the business cycle turning points identified are very similar across approaches, we show that SARB's Leading Indicator does not substantively improve predictions of the business cycle relative to GDP itself. We therefore assess the performance of a range of potential leading indicators in identifying downturns and consider whether alternative indicators and estimation approaches can produce better predictions than those of the SARB. We demonstrate that using a larger information set produces substantially better business cycle predictions, especially when using machine learning techniques.

[1] Cycle durations are constrained so that peak-to-trough and trough-to-peak phases last at least two quarters, while a full cycle (peak-to-peak or trough-to-trough) must span a minimum of five quarters - ensuring that short-term fluctuations, such as a single quarter's movement in the terms of trade, are not classified as cycles. These duration thresholds align with those used in cross-country analyses of GDP cycles. The BBQ algorithm is applied to the cyclical component of seasonally adjusted real GDP obtained using a Hodrick-Prescott filter (Hodrick and Prescott, 1997) with a lambda of 16 000 for quarterly data, with prior adjustment using an Hodrick-Prescott filter with lambda of 500 to replace extreme outliers such as those associated with the COVID-19 pandemic of 2020. This is similar to the approach used by SARB, although SARB apply the filter to monthly time series with a lamdha of 108 000, see Venter and Wolhuter (2023).

Figure 1: Illustration of the South African Business Cycle

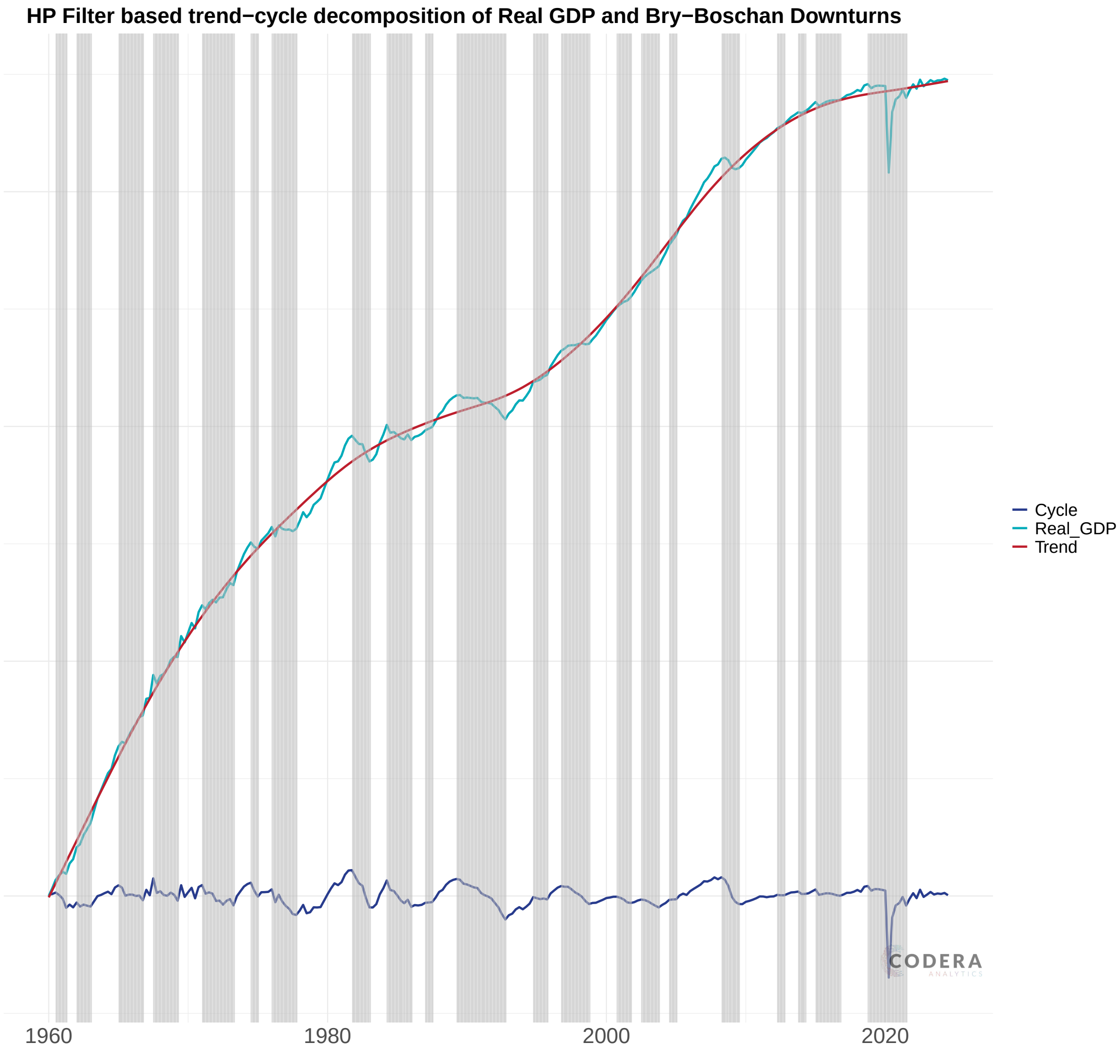

2 Identifying South African business cycles

Figure 2 compares business cycle phases identified using the BBQ algorith and those identified by the SARB and OECD. The figure shows that the three approaches identify broadly similar turning points. The BBQ algorithm detects more frequent business cycle turning points, with slightly shorter cycle durations compared to those identified by the OECD and SARB. The SARB turning points generally align with those of the OECD, though SARB does not recognise one turning point identified by both the OECD as a business cycle turning point. Since 1995, BBQ business algorithm identifies 10 cycles with 10 downward phases, having an average during of 6 quarters. The downward phases have had an average duration of 5 quarters and the expansion phases have a average duration of 7 quarters. The SARB business cycles have had 4 cycles with 4 downward phases, having an average duration of 13 quarters. The downward phases had an average duration of 9 quarters and the expansion phases had an average duration of 17 quarters over the period. The OECD business cycles have had 5 cycles with 5 downward phases, having an average duration of 10 quarters. The downward phases have an average duration of 9 quarters and the expansion phases have an average duration of 11 quarters.

Figure 2: Comparison of identified South African Business Cycles (1=Downturn)

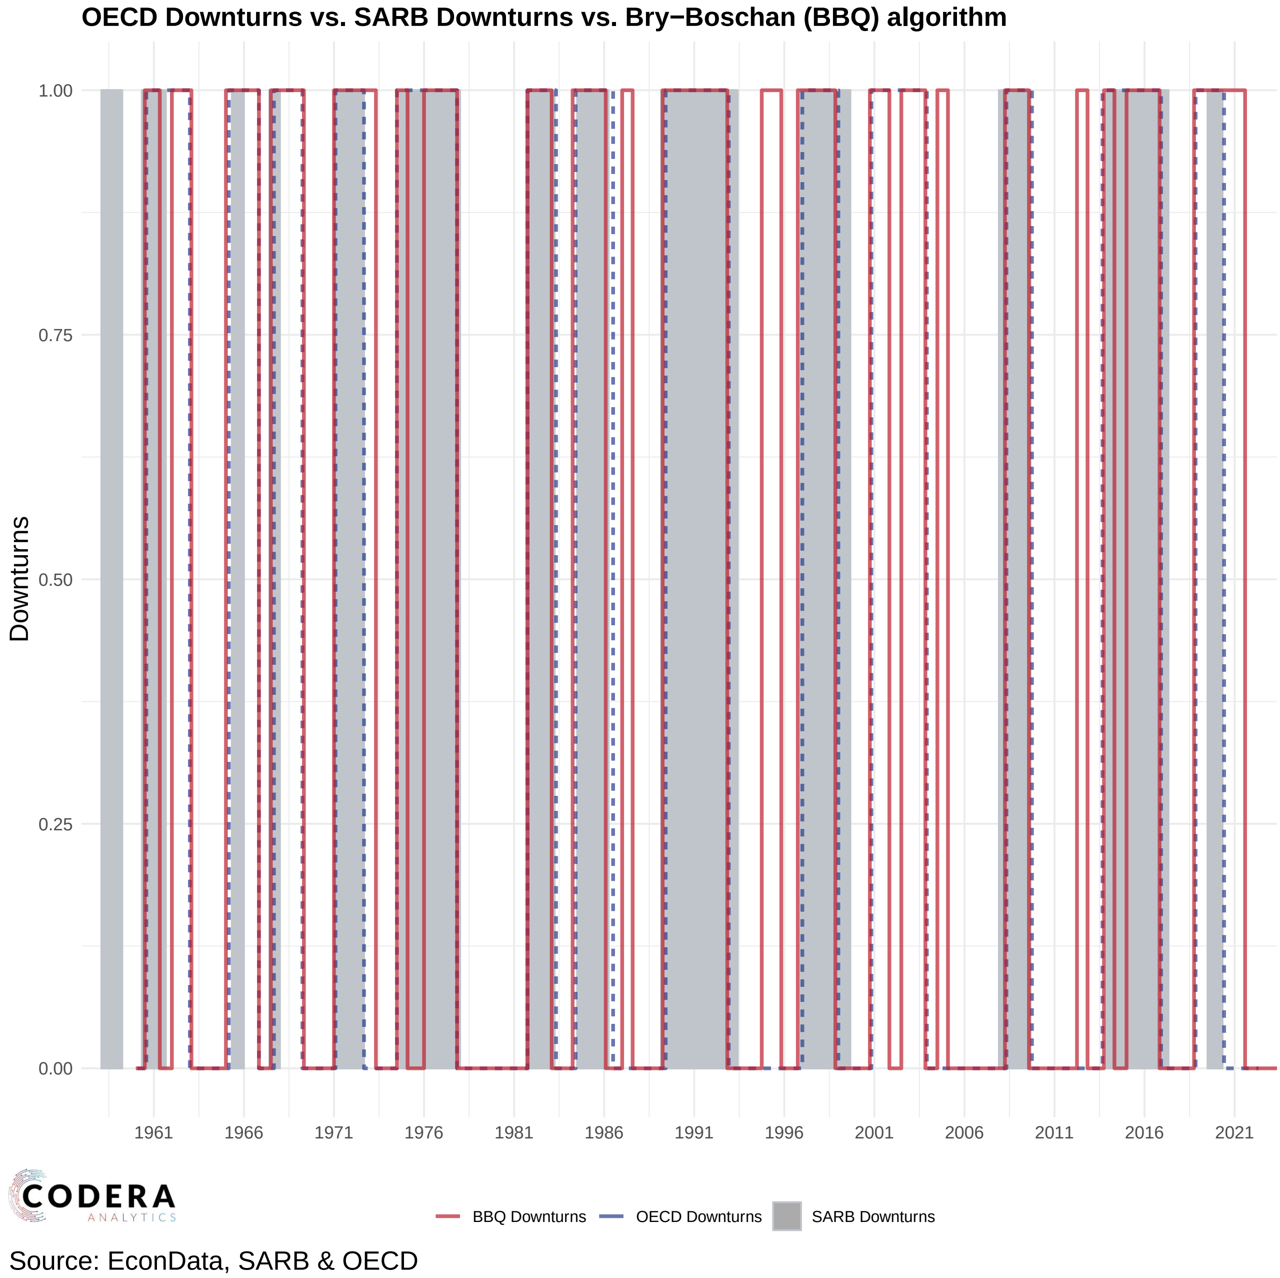

A dearth of recent papers estimating business cycle phases across countries makes it difficult to easily compare our estimates to those of other countries. Generally, empirical papers find that contraction phases tend to be shorter than expansion phases in both advanced and emerging market countries. What makes South Africa's business cycles unusual compared to peer emerging market economies is how long downswings have been, how short upswings have been and the low amplitude of expansions. Calderon and Fuentes (2010) show that this was already the case before the global financial crisis, but Table 1 shows that there a substantial decline in the duration and amplitude of expansions in South Africa post-global financial crisis (GFC). There have not been enough cycles post-COVID-19 pandemic to draw comparisons to, though annual GDP growth has been dramatically lower than before the GFC. Downward cycles have also become larger and longer post-GFC. The consequence of this has been a stagnation in the level of real GDP over the last decade, and a decline in real GDP per capita. Had the economy continued to grow at its' long term trend growth rate pre-GFC/COVID, per capita real GDP would have been at least 14%/12% higher currently. If the growth rate between 2000 and the GFC could have been sustained, it would have been more than 30% higher in just 15 years.[2] Lastly, these estimates demonstrate that South Africa has a very volatile business cycle, making it particularly difficult to predict turning points in the economic cycle.

Downward Phase	**1994Q1–2008Q1**	**2010Q1–2019Q4**	**2020Q4–2024Q4**
Mean duration (quarters)	5	6	–
Mean amplitude (%)	2.11	5.83	–
Upward Phase	**1994Q1–2008Q1**	**2010Q1–2019Q4**	**2020Q4–2024Q4**
Mean duration (quarters)	6	5	14
Mean amplitude (%)	5.11	1.54	2.53

Table 1: Characterising business cycles across sub-samples (BBQ cycles)

3 Predicting Downward Business Cycle Phases

The aim of this paper is to assess the ability of different potential indicators to predict downward phases of the South African business cycle. Any useful leading indicator needs to be better than GDP itself at predicting future business cycle turning points. This is the purpose of the SARB Leading Indicator: to provide early signals of turning points in the business cycle (SARB, 2024) so we start by assessing the SARB's indicator's ability to predict business cycle turning points. Since the OECD estimates for South Africa are not available beyond 2022 and SARB's estimates are updated only quarterly, we use BBQ business cycles as our benchmark.

In this section, we assess whether different economic and financial indicators serve as good predictors of the South African busincess cycle.

[2] See this Codera blog post for more detail. Also see this Codera post for a cross-country comparison.

3.1 Dataset

We consider a large number of potential leading indicators, including those used by Botha et al. (2021), and as the methodological section that follows shows, we also consider aspects of their statistical properties. We exclude series that are not available at least from 2000 onwards, or are not regularly publicly available so that indicator creation can be automated.[3] We focus on ex-post analysis (i.e. using revised data as currently available).[4] SARB's composite leading business cycle indicator is based on eleven indicators, while the OECD measure includes six.[5] Unfortunately, not all of the these components are publicly available on an ongoing basis so we do not replicate SARB's indicator to provide real-time estimates, instead relying on ex-post (i.e. revised estimates) as published by SARB.

Figure 3 shows that confidence indicators are weakly contemporaneously correlated with business cycle phases, but not highly correlated with GDP growth outcomes in South Africa. SARB's coincident and leading indicators are not highly correlated with the business cycle or other cyclical macroeconomic dynamics.[6]

[3]Another potential indicator worth mentioning is BankServ Africa's banking activity index (BETI) as it is available with a short lead time, making it a potentially useful source of information with which to nowcast GDP. It is not publicly available for the full sample we consider, but its contemporaneous correlation to GDP is around 0.6, 0.3 one-quarter ahead and 0.1 two quarters-ahead. Visit this Codera blog post or this one for more information. Future policy notes will consider unbalanced datasets to include indicators that are not available since 2000. Note that we backfilled two quarterly observations for inflation expectations and three for retail trade to include these in a balanced dataset. This did not have a qualitative impact on the empirical results.

[4]Our EconData platform provides truly 'real-time' data (i.e. unrevised data as available in a given month). Contact us for a demonstration of the performance of our real-time indicators. We plan to publish real-time leading indicators in future policy briefs.

[5]These are job advertisement space in the Sunday Times newspaper, Number of residential building plans passed for flats, townhouses and houses larger than 80 m squared, Interest rate spread, Real M1 money supply (deflated with CPI): six-month smoothed growth rate, Index of commodity prices (in US dollar) for a basket of South African-produced export commodities, Composite leading business cycle indicator of South Africa's major trading partner countries: percentage change over twelve months, Gross operating surplus as a percentage of gross domestic product, RMB/BER Business Confidence Index, Net balance of manufacturers observing an increase in the average number of hours worked per factory worker (half weight), Net balance of manufacturers observing an increase in the volume of domestic orders received (half weight), Number of new passenger vehicles sold: percentage change over twelve months. The OECD leading indicator includes Manufacturing Orders, Manufacturing Confidence, Permits Issued for Total Buildings, Passenger Car Registrations, FTSE/JSE Index, Spread of Interest Rates. SARB also uses diffusion indices that track the proportion of indicator time series rising relative to their trends to confirm turning points in the business cycle, see Venter and Wolhuter (2023).

[6]This is consistent with this earlier Codera blog post that compares leading indicator correlations with GDP growth outcomes in South Africa.

Table 2: Data Series Considered

Name	**Description**	**Source**
Real GDP	Seasonally Adjusted (2015 base) in Rand (Millions)	SARB, EconData
Consumer Confidence Index	FNB/BER Index	Bureau for Economic Research (BER)
Business Confidence Index	RMB/BER Index	BER
SACCI Business Confidence Index	Index	South African Chamber of Commerce and Industry (SACCI)
Manufacturing PMI	Seasonally Adjusted Monthly Survey Index	Absa
Consumer Price Inflation	Index	Statistics South Africa (Stats SA), EconData
3-Month Gov Bond Yield	Percentage Terms	OECD
10-Year Gov Bond Yield	Percentage Terms	OECD
Inflation Expectations: One Year Ahead	BER Index	SARB, BER, EconData
Inflation Expectations: Analyst Current Year	Not Seasonally Adjusted in Percentage Terms	SARB, BER, EconData
JSE All Share Index (monthly mean)	Mean of the Index's daily closing values over a month	WSJ, FactSet
JSE All Share Index (monthly standard deviation)	Standard Deviation of the Index's daily closing values over a month	WSJ, FactSet
SARB Leading Indicator	Seasonally Adjusted (2015 base) Index	SARB, EconData
SARB Coincident Indicator	Seasonally Adjusted (2015 base) Index	SARB, EconData
Leading Indicator USA	SARB Index	SARB, EconData
Leading Indicator Trade Partners Excl USA	SARB Index	SARB, EconData
Leading Indicator All Trade Partners	SARB Index	SARB, EconData
Coincident Indicator All Trading Partners	SARB Index	SARB, EconData
Total Earnings	Remuneration per worker in non-agricultural sectors. Current prices. Seasonally Adjusted (2010 base) Index	SARB, EconData
Unemployment Rate	Seasonally Adjusted Percentage	SARB, EconData
Manufacturing Production	Volumes and Values, Seasonally Adjusted (2015 base)	SARB, EconData
Producer Prices	Seasonally Adjusted (December 2010 base) Index	SARB, EconData
Consumption Expenditure Households	Seasonally Adjusted (2015 base) in Rand (Millions)	SARB, EconData
Consumption Expenditure General Government	Seasonally Adjusted (2015 base) in Rand (Millions)	SARB, EconData
Gross Fixed Capital Formation	Seasonally Adjusted (2015 base) in Rand (Millions)	SARB, EconData
Total Mining Production	Seasonally Adjusted (2015 base) Index	SARB, EconData
Electricity Generated	Seasonally Adjusted (2015 base) Index	SARB, EconData
Delivery Period of Orders Received in Manufacturing	Values	BER,SARB, EconData
Total Retail Trade	Seasonally Adjusted (2015 base) in Rand (Millions)	Stats SA, EconData
Total Liquidations	Neither Seasonally nor Calendar Adjusted as Total Number	Stats SA, EconData
Mining Volume	Seasonally Adjusted (2019 base) Index	Stats SA, EconData
Electricity Volume	Seasonally Adjusted (2015 base) Index	Stats SA, EconData
Manufacturing Volume	Seasonally Adjusted (2019 base) Index	Stats SA, EconData
Exports of Goods and Services	Seasonally Adjusted (2015 base) in Rand (Millions)	SARB, EconData
Imports of Goods and Services	Seasonally Adjusted (2015 base) in Rand (Millions)	SARB, EconData
Total Motor Trade	Seasonally Adjusted in Rand (Millions)	Stats SA, EconData
Purchasing Managers Index: Prices	Not Seasonally or Calendar Adjusted Index	SARB, EconData

Name	Description	Source
Rand to Dollar Exchange Rate	Monthly Average Rate in Rand	SARB, EconData
Rand to Pound Exchange Rate	Monthly Average Rate in Rand	SARB, EconData
Rand to Euro Exchange Rate	Monthly Average Rate in Rand	SARB, EconData
Rand to Yen Exchange Rate	Monthly Average Rate in Rand	SARB, EconData
Rand	Gold Price Quoted in Rand	SARB, EconData
Geopolitical Risk Index (Global)	Index	Caldara and Iacoviello (2022)
Geopolitical Risk Index (South Africa)	Index	Caldara and Iacoviello (2022)
Yield Curve Spread	Percentage difference between 3 and 10 Month Yield	Derived using data from OECD
Downward Phase Indicator	Binary Indicator for Business Cycles (0 or 1)	SARB & OECD
Number of Credit Card Purchases	Not Seasonally or Calendar Adjusted Total (Millions)	SARB, EconData
Value of Credit Card Purchases	Seasonally Adjusted Value in Rand (Millions)	SARB, EconData
Electronic Funds Transfers	Seasonally Adjusted Value in Rand (Millions)	SARB, EconData
Credit Extended to Private Sector	Seasonally Adjusted Value in Rand (Millions)	SARB, EconData

Figure 3: Correlation of selected leading indicators and South African Business Cycle

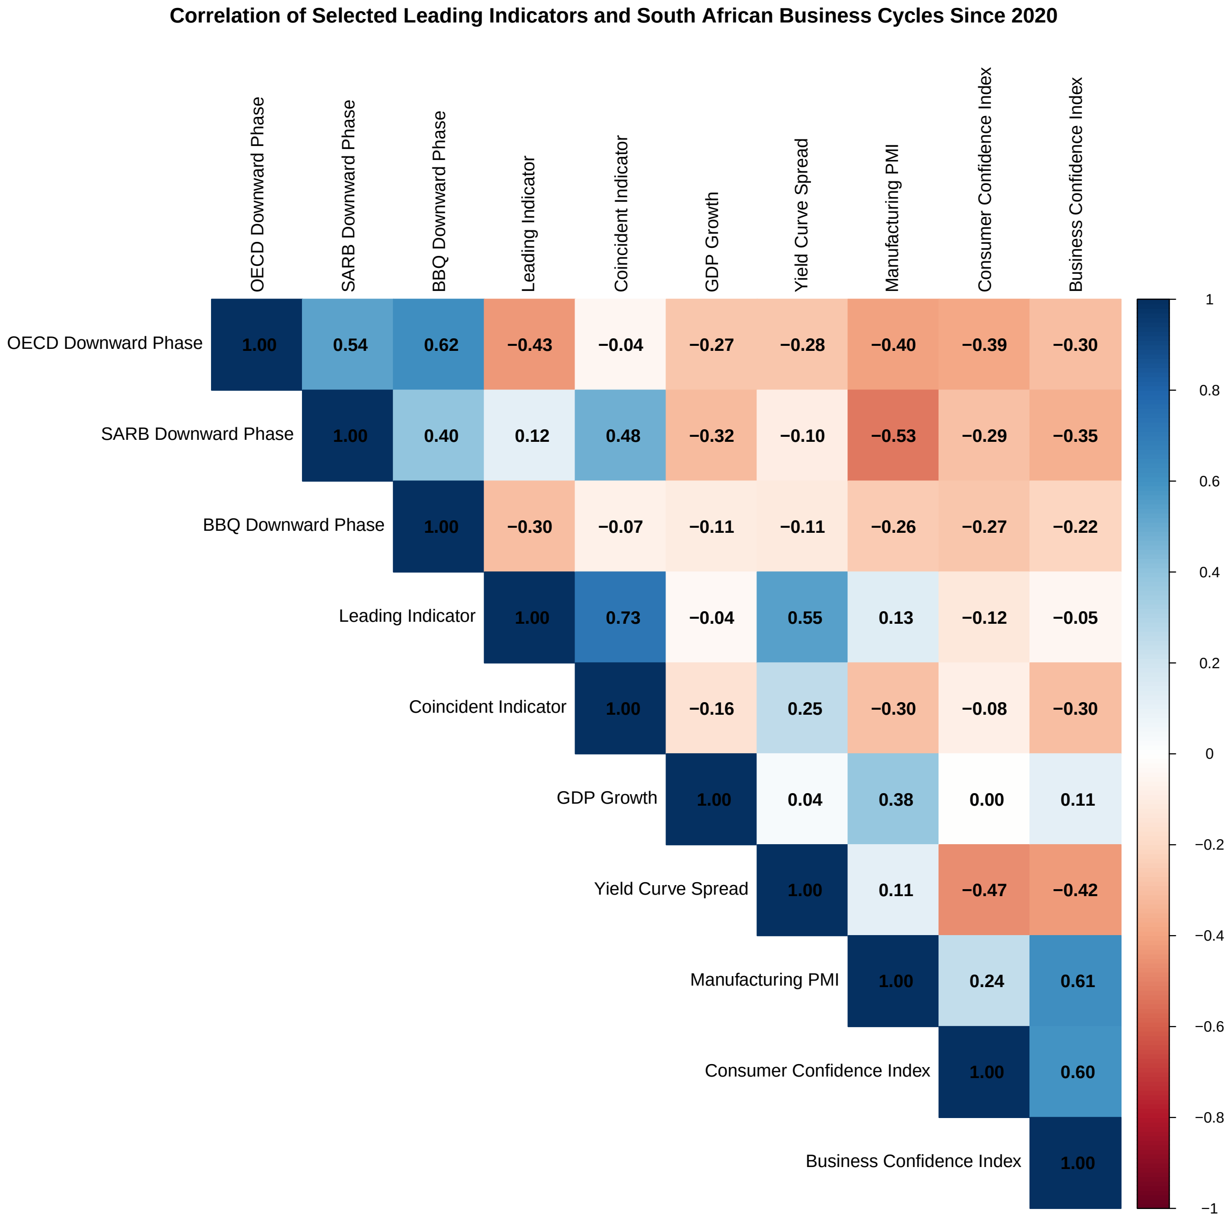

Note: Pairwise correlations between series are plotted over matching samples at quarterly frequency.

3.2 Modelling business cycle predictors

We take a 'kitchen-sink' approach to assessing the performance of different indicators in predicting business cycles, considering not only a large number of potential predictors, but also their statistical properties and a variety of simple and complex forecasting approaches. We consider Logit, Probit, and Least Absolute Shrinkage and Selection Operator (LASSO) approaches as these frameworks are good at binary classification and probability estimation and can accommodate large datasets. We also consider machine learning (Random Forest) model specifications that enable a large dataset to be used and complex relationships to be modelled. Our baseline approach is simple Probit models using real GDP as the predictor with some simple transformations done as presented below.

The features used in the Random Forest, Logit and Probit models for the baseline GDP models over different time horizons are:

- **Change over given horizon:** Change in Real GDP over 1, 2 and 4 quarter horizons.
- **Rolling mean over the horizon:** The rolling mean of Real GDP over 1, 2 and 4 quarter horizons.

For the leading indicator models, we use feature transformations to represent the data at a quarterly frequency. The features used in the Random Forest, Logit and Probit models for the leading indicator models over different time horizons are:

- **Change over given horizon:** Difference in the indicator over 1-quarter (3-month), 2-quarter (6-month), and 4-quarter (12-month) horizons.
- **Rolling mean over the horizon:** Right-aligned moving average of the indicator over 1-quarter (3-month), 2-quarter (6-month), and 4-quarter (12-month) horizons.

For models that include the broader set of features, any monthly series are transformed to a quarterly frequency using right-aligned rolling averages over the relevant 3 month period. All model features are standardised to have zero mean and unit variance to ensure that no single variable dominates the model owing to its scale. This also supports faster andmore stable convergence during training for the Logit and Probit models.

All models used in this paper were estimated using a rolling window approach, where models are trained on a fixed-length subset of the data and then used to generate a forecast for the next period. After each prediction, the window is moved forward by one observation, the earliest point is dropped, and the newest data point is added. This approach simulates real-time forecasting and helps evaluate how model performance evolves over time as new information becomes available.

To estimate Lasso regression models, we use a rolling window forecast with a 60/40 training-test split. Within each training window, we estimate a penalized logistic regression model using a Lasso penalty. The Lasso penalty encourages sparsity in the coefficient estimates, with the shrinkage parameter selected using k-fold cross-validation to select the parameter value that minimises prediction error. The fitted model is then used to generate out of sample probabilitiy forecasts for the next observation.

For Random Forest models, we implemented a 60/40 training-test split, chosen to ensure at least two downward phases were available for validation, and a rolling window forecast approach. We use a grid search with time series cross-validation to determine the optimal number of predictors that will be randomly sampled at each split when creating the tree models, followed by a tuning loop to select the best number of trees based on prediction accuracy. With these hyper-parameters in hand, models are retrained for each new observation using all prior data, ensuring no data leakage. In parallel, logistic regression models (both probit, logit and penalized logit) are implemented within the same rolling framework and cross validated to predict downward phase probabilities. The coefficient magnitudes served as indicators of feature importance over time (for the penalized models), and additional diagnostic checks - such as Variance Inflation Factors (VIF) for multicollinearity and the Hosmer-Lemeshow test for goodness-of-fit - were performed for select windows to ensure model reliability.[7] Further details regarding the methodologies applied are available in Appendix A.

[7]Note that no pre-processing was done to the features used in the Random Forecast model. Since many features are highly correlated, in future work we will consider additional frameworks for dimension reduction, use of mixed frequencies and for isolating drivers of business cycle phases.

4 Ability to predict Business Cycles

A good leading indicator should be able to predict business cycle turning points. Table 3 explains the metrics used to assess model accuracy in business phase prediction.

Metric	Description
Accuracy	The proportion of all predictions that were correct.
Balanced Accuracy	The average of sensitivity and specificity, adjusting for class imbalance.
Specificity	The proportion of actual negatives correctly identified as negative.
Sensitivity	The proportion of actual positives correctly identified as positive.
F1-Score	The harmonic mean of precision and sensitivity, balancing false positives and false negatives.

Table 3: Classification metrics used to evaluate model performance.

We begin with a comparison of models and indicators in predicting BBQ-based downward cycles, focusing on just models that only incorporate GDP or Leading Indicator, respectively and the Random Forest model drawing on the full set of indicators at a one-quarter ahead forecast horizon. Results for benchmark GDP models, leading indicator models and various model frameworks drawing on the full data set across one-, two-, and four-quarters ahead forecasting BBQ, SARB and OECD business cycles are available in Appendix B.

Table 4 shows that the SARB Leading Indicator produces forecasts of the business cycle are similar in accuracy as just using GDP itself for BBQ cycles.[8] We observe clear improvements in model performance when including a large dataset of indicators. The Random Forest model provides substantial accuracy gains across all metrics, using a broader feature set than both the benchmark GDP Probit and the SARB Leading Indicator Probit models. This improvement is achieved without any explicit feature engineering, highlighting the value in the more expanded feature set. The Random Forest performance can be attributed to its ability to capture non-linear relationships and complex interactions between predictors automatically.

Sensitivity and specificity provide measures of how well a model identifies downward phases identified by BBQ, and how well the model identifies upward phases identified by BBQ, respectively. The GDP Probit and SARB Probit models tend to be skewed toward specificity, with values of 0.75 and 0.84 respectively, suggesting a tendency to avoid false positives in identifying business cycles. However, this comes at the cost of lower sensitivity - 0.47 for the GDP Probit and 0.6 for the SARB indicator model - resulting in more missed downward phases. In contrast, the Random Forest models demonstrate a more balanced trade-off between sensitivity and specificity, achieving a sensitivity of 0.85 and a specificity of 0.79, reflecting a strong ability to correctly classify both expansion and downward phases. This improved trade-off is reflected in the higher F1-score of 0.83, indicating stronger overall classification performance when both

[8]We show in a different context that the SARB Leading Indicator does not produce better forecasts than a simple AR(1) model (that just uses past values of GDP to project future values) at one-quarter-ahead and one-year-ahead horizons, see this Codera blog post. In a follow-up paper, we plan to publish a leading indicator optimised for business cycle identification and GDP nowcasting.

false alarms and missed signals are costly - an important consideration in the context of business cycle analysis.

In Tables 5 and 6 we repeat the excercise, comparing the ability of GDP alone, the SARB Leading Indicator alone, and Random Forest models that incorporate a large dataset of indicators to predict downturns identified by the SARB and OECD, respectively. The results are qualitatively comparable to those from Table 4, but with less accurate predictions for SARB cycles from the SARB Leading Indicator. Models that include a larger information set achieve better prediction accuracy, with the Random Forest model the best performing framework across all datasets and business cycle chronologies.

To illustrate how these models can be used to identify potential indicators of business cycles, we plot the top predictors of BBQ, SARB and OECD business cycle downturns from the Random Forest model from these tables (Figures 4, 5 and 6). The exchange rate is the most important predictor for cycle downturns in the OECD and SARB Random Forest models, with the bond yield being the most important for the BBQ model. The other important indicators are coincident indicators, gross fixed capital formation, and average wages. Although the SARB Leading Indicator makes an apperance in the top 10 predictors for the OECD cycle predictions, it is not a very good predictor of cycles by itself. Across these models, the top indicators straddle real-activity measures, price-pressure measures, and financial-condition metrics, supporting the case for using a larger information set when predicting business cycle dynamics, as well as constructing indicators using an equally-weighted index.

	GDP Probit (0.5)	**SARB Leading Indicator Probit (0.35)**	**Indicator Random Forest (0.6)**
Accuracy	0.6	0.72	0.82
Balanced Accuracy	0.61	0.72	0.82
Sensitivity	0.47	0.6	0.85
Specificity	0.75	0.84	0.79
F1-Score	0.56	0.69	0.83

Table 4: Performance Metrics for BBQ Phases (one-quarter ahead)

	GDP Probit (0.65)	**SARB Leading Indicator Probit (0.2)**	**Indicator Random Forest (0.5)**
Accuracy	0.64	0.56	0.85
Balanced Accuracy	0.62	0.54	0.85
Sensitivity	0.7	0.72	0.83
Specificity	0.53	0.36	0.87
F1-Score	0.72	0.65	0.86

Table 5: Performance Metrics for SARB Business Cycle Phases (one-quarter ahead)

	GDP Probit (0.45)	**SARB Leading Indicator Probit (0.3)**	**Indicator Random Forest (0.55)**
Accuracy	0.61	0.69	0.83
Balanced Accuracy	0.70	0.73	0.83
Sensitivity	0.63	0.53	0.84
Specificity	0.59	0.88	0.82
F1-Score	0.63	0.65	0.84

Table 6: Performance Metrics for OECD Business Cycle Phases (one-quarter ahead)

Note: The values in brackets represent the selected threshold for the model above which a downward phase is predicted. For example, if the threshold is 0.5, then a probability of 0.6 would be classified as a downward phase, while a probability of 0.4 would not be classified as a downward phase.

Figure 4: Most important indicators for BBQ downturn prediction (Random Forest)

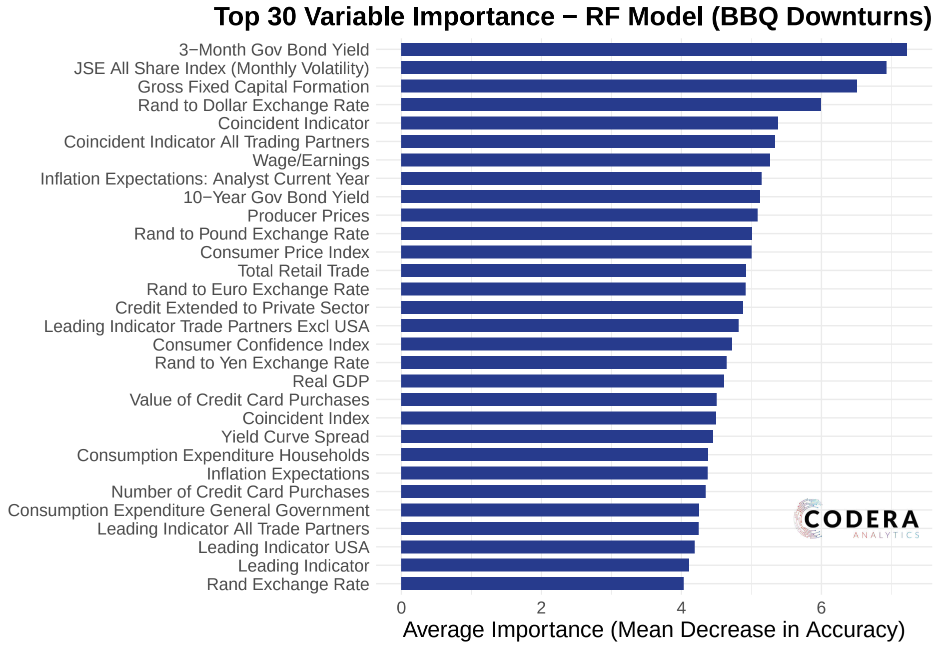

Figure 5: Most important indicators for SARB downturn prediction (Random Forest)

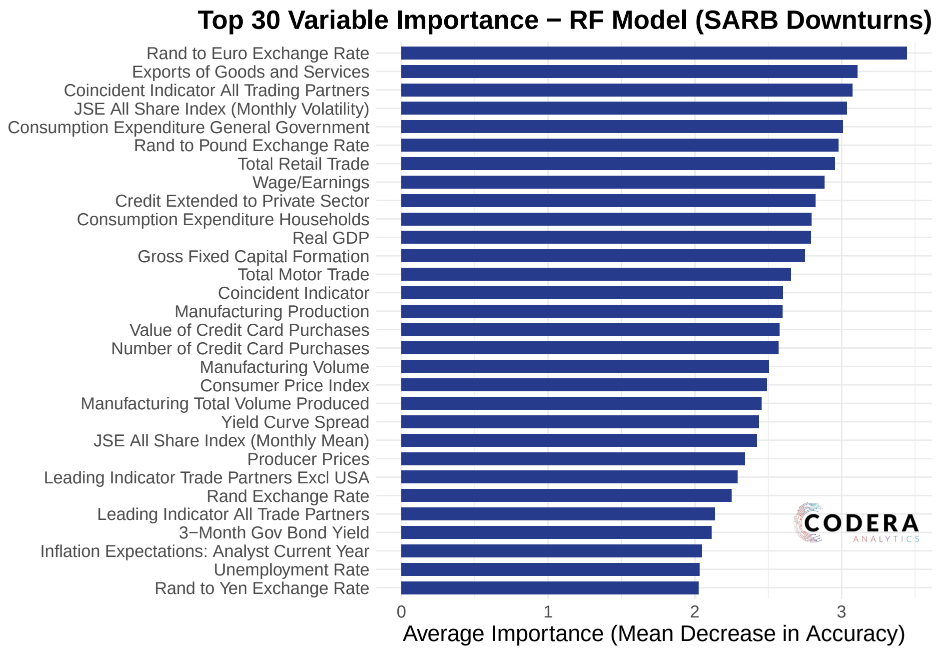

Figure 6: Most important indicators for OECD downturn prediction (Random Forest)

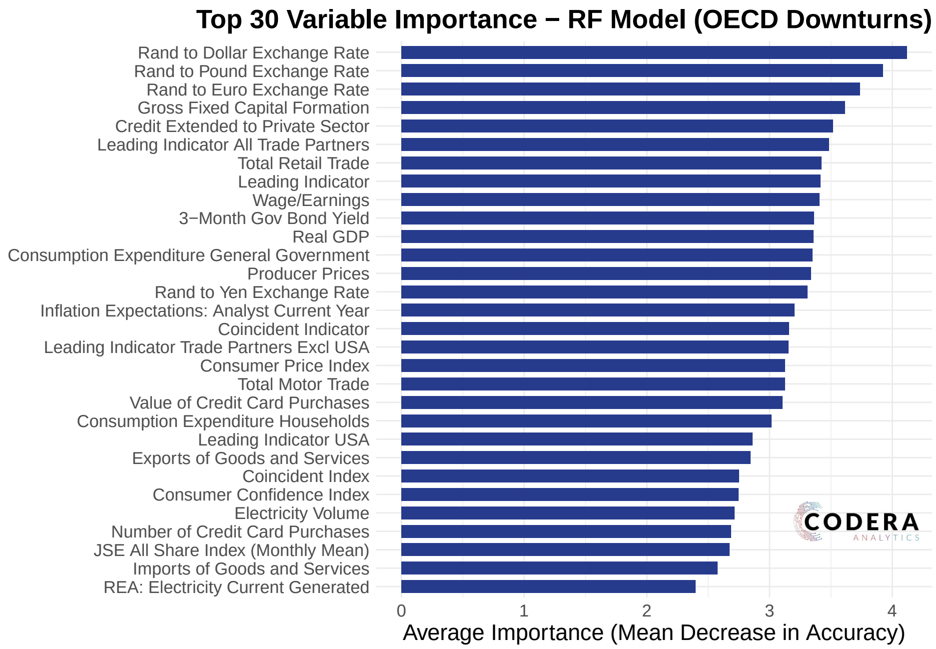

5 Conclusion

This paper has demonstrated that South Africa has a very volatile business cycle that makes it particularly difficult to predict turning points in the economic cycle. South Africa's business cycle is characterised by relatively long downswings and short upswing phases with low amplitude. The consequence of this has been a stagnation in the level of real GDP over the last decade, and a decline in real GDP per capita.

We demonstrate that the SARB's Leading Indicator does not substantively improve predictions of the business cycle relative to GDP itself. We show that using a large number of indicators produces substantially better business cycle predictions, especially when using machine learning techniques. These results underscore the value of using a large information set when predicting business cycle dynamics. Our findings also have implications for the creation of composite leading indicators, with our results suggesting that many of the macroeconomic variables considered by analysts as leading indicators do not provide good signals of GDP growth or developments in the South African business cycle. There are many possible extensions to the analysis presented, including using a true real-time dataset and creating indicators for tracking growth in real time, using big data and mixed frequency machine learning approaches to nowcast business cycle dynamics as well as GDP, and assessing the implications of business cycle movements for the structural changes in the economy.

References

Botha, B., Olds, T., Reid, G., Steenkamp, D., and van Jaarsveld, R. (2021). Nowcasting South African gross domestic product using a suite of statistical models. *South African Journal of Economics*, 89(4):526–554.

Bry, G. and Boschan, C. (1971). *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs*. National Bureau of Economic Research, Inc.

Caldara, D. and Iacoviello, M. (2022). Measuring Geopolitical Risk. *American Economic Review*, 112(4):1194–1225.

Calderon, C. and Fuentes, R. (2010). Characterizing the business cycles of emerging economies. Policy Research Working Paper Series 5343, The World Bank.

Harding, D. and Pagan, A. (2002). Dissecting the cycle: a methodological investigation. *Journal of Monetary Economics*, 49(2):365–381.

Hodrick, R. and Prescott, E. (1997). Postwar u.s. business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, 29(1):1–16.

SARB (2024). Quarterly Bulletion Box 1: Revisions to the composite leading and coincident business cycle indicators. South African Reserve Bank Quarterly Bulletin December 2024.

Venter, J. C. (2020). *Assessing the 2013 and 2017 Business Cycle Turning Points Signalled by the SARB's Composite Leading Business Cycle Indicator*, pages 265–284. Springer International Publishing, Cham.

Venter, J. C. and Pretorius, W. S. (2004). Note on the revision of composite leading and coincident business cycle indicators. South African Reserve Bank Quarterly Bulletin March 2004.

Venter, J. C. and Wolhuter, A. (2023). The South African business cycle from 2013 to 2022. South African Reserve Bank Quarterly Bulletin March 2023.

Appendices

Appendix A Training/Test Split

We employ a rolling window framework to evaluate the predictive performance of our model while ensuring that only past information is used for forecasting. Initially, the first 60% of the chronologically ordered dataset is designated as the training set. For example, if N is the number of observations, the initial window is set to 0.6N observations. The model is first trained on this initial subset. Thereafter, for each subsequent time step t, the model is re-estimated using all available data from time 1 up to t-1 (i.e., the training set is expanded incrementally), and a prediction is made for the observation at time t. This sequential approach mimics a real-time scenario and eliminates look-ahead bias, as no future data are used during model estimation. The performance of the model is then assessed on the out-of-sample predictions generated in this manner.

A.0.1 Logit and Probit Regressions

For the benchmark model where we use real GDP as our predictor, we consider the following features:

- **real GDP change**: 1, 2 or 4-quarter change in real GDP.
- **real GDP trend**: 1, 2 or 4-quarter trend component of the real GDP.

For each time step in the test period (i.e., after the initial 60% training data), the model is re-estimated using historical data. In each iteration, we extract all observations up to time t-1 and compute the relevant features.

We use a time-series cross-validation within each rolling window, with an initial window set to 80% of the available training set (or at least 10 observations) and a forecast horizon of one-time step. This ensures that the model is tuned using only past data, avoiding data leakage. We fit a logistic regression model on the training data. The fitted model is used to predict the probability of a downward phase for the test observation at time t. For select windows we perform diagnostic checks - computing Variance Inflation Factors (VIF) to assess multicollinearity and applying the Hosmer-Lemeshow test to evaluate goodness-of-fit. This rolling window approach guarantees that each forecast is made using only information available up to the time of prediction, while the time-series cross-validation scheme helps in obtaining robust estimates of model performance.

A.0.2 Random Forest model estimation

For the benchmark model where we use real GDP as our predictor, we consider the following features:

- **real GDP change**: 1, 2 or 4-quarter change in real GDP.
- **real GDP trend**: 1, 2 or 4-quarter trend component of the real GDP.

The evaluation approach for the Random Forest method involves both hyper-parameter tuning and a rolling window forecast strategy to mimic a realistic time-series forecasting scenario. Initially, 60% of the data is used as the training set, where features are generated without data leakage. A grid search is done for the number of features (predictors) randomly sampled as candidates at each split in a Random Forest using time-series cross-validation with a timeslice method - this sets up a fixed window with a horizon of one observation, ensuring that the temporal order is respected. After identifying the best number of features randomly sampled, a range of number of trees in the forest are tested, and the model with the highest accuracy is selected.

Once optimal hyper-parameters have been determined, the model is evaluated through a rolling window approach. For each time point in the test set, all preceding data is used to train the model, and features are recomputed to maintain integrity. The trained model then produces predicted probabilities for the next observation, which are thresholded to generate binary predictions. These predictions are compared against actual downward phase indicators using a confusion matrix, from which metrics such as accuracy, balanced accuracy, sensitivity, specificity, and F1-score are derived. This evaluation strategy assesses in-sample performance and also simulates out-of-sample forecasting, providing a robust measure of the model's real-world predictive ability.

Appendix B Additional results

B.1 Benchmark GDP model results

B.1.1 BBQ Business Cycle Phase

	1Q Ahead (0.5)	**2Q Ahead (0.5)**	**4Q Ahead (0.4)**
Accuracy	0.6	0.53	0.41
Balanced Accuracy	0.61	0.54	0.41
Sensitivity	0.47	0.33	0.63
Specificity	0.75	0.75	0.19
F1-Score	0.56	0.43	0.52

Table 7: Performance Metrics for Bry-Boschan Business Cycle Phase: Benchmark Probit

	1Q Ahead (0.55)	**2Q Ahead (0.55)**	**4Q Ahead (0.5)**
Accuracy	0.6	0.53	0.41
Balanced Accuracy	0.61	0.54	0.41
Sensitivity	0.47	0.33	0.19
Specificity	0.75	0.75	0.63
F1-Score	0.56	0.43	0.52

Table 8: Performance Metrics for Bry-Boschan Business Cycle Phase: Benchmark Logit

	1Q Ahead (0.5)	**2Q Ahead 0.55**	**4Q Ahead (0.7)**
Accuracy	0.71	0.68	0.73
Balanced Accuracy	0.71	0.67	0.73
Sensitivity	0.68	0.72	0.83
Specificity	0.75	0.63	0.63
F1-Score	0.72	0.7	0.75

Table 9: Performance Metrics for Bry-Boschan Business Cycle Phase: Benchmark Random Forest

B.1.2 SARB Business Cycle Phase

	1Q Ahead (0.65)	**2Q Ahead (0.6)**	**4Q Ahead (0.55)**
Accuracy	0.64	0.63	0.53
Balanced Accuracy	0.62	0.58	0.45
Sensitivity	0.7	0.75	0.7
Specificity	0.54	0.42	0.2
F1-Score	0.7	0.71	0.67

Table 10: Performance Metrics for SARB Business Cycle Phase: Benchmark Logit

	1Q Ahead (0.65)	**2Q Ahead (0.65)**	**4Q Ahead (0.6)**
Accuracy	0.64	0.59	0.43
Balanced Accuracy	0.62	0.56	0.38
Sensitivity	0.7	0.7	0.55
Specificity	0.53	0.41	0.2
F1-Score	0.72	0.68	0.56

Table 11: Performance Metrics for SARB Business Cycle Phase: Benchmark Probit

	1Q Ahead (0.4)	2Q Ahead (0.5)	4Q Ahead (0.5)
Accuracy	0.82	0.75	0.8
Balanced Accuracy	0.79	0.75	0.8
Sensitivity	0.9	0.75	0.8
Specificity	0.69	0.75	0.8
F1-Score	0.86	0.79	0.84

Table 12: Performance Metrics for SARB Business Cycle Phase: Benchmark Random Forest

B.1.3 OECD Business Cycle Phase

	1Q Ahead (0.45)	2Q Ahead (0.5)	4Q Ahead (0.5)
Accuracy	0.61	0.64	0.47
Balanced Accuracy	0.61	0.64	0.48
Sensitivity	0.63	0.68	0.26
Specificity	0.59	0.59	0.71
F1-Score	0.63	0.67	0.34

Table 13: Performance Metrics for OECD Business Cycle Phase: Benchmark Probit

	1Q Ahead (0.5)	2Q Ahead (0.6)	4Q Ahead (0.5)
Accuracy	0.75	0.72	0.5
Balanced Accuracy	0.75	0.72	0.51
Sensitivity	0.79	0.68	0.32
Specificity	0.71	0.76	0.71
F1-Score	0.77	0.72	0.4

Table 14: Performance Metrics for OECD Business Cycle Phase: Benchmark Logit

	1Q Ahead (0.5)	2Q Ahead (0.5)	4Q Ahead (0.65)
Accuracy	0.75	0.72	0.78
Balanced Accuracy	0.76	0.72	0.78
Sensitivity	0.79	0.79	0.79
Specificity	0.72	0.65	0.76
F1-Score	0.73	0.75	0.79

Table 15: Performance Metrics for OECD Business Cycle Phase: Benchmark Random Forest

B.2 SARB Leading Indicator Model Results

B.2.1 BBQ Business Cycle Phase

	3 Months Ahead (0.35)	6 Months Ahead (0.35)	12 Months Ahead (0.4)
Accuracy	0.72	0.74	0.68
Balanced Accuracy	0.72	0.74	0.69
Sensitivity	0.6	0.85	0.89
Specificity	0.84	0.63	0.5
F1-Score	0.69	0.77	0.73

Table 16: Performance Metrics for Bry-Boschan Business Cycle Phase: Leading Indicator Logit

	3 Months Ahead (0.4)	6 Months Ahead (0.4)	12 Months Ahead (0.4)
Accuracy	0.74	0.74	0.71
Balanced Accuracy	0.75	0.74	0.72
Sensitivity	0.6	0.7	0.89
Specificity	0.79	0.78	0.55
F1-Score	0.71	0.74	0.74

Table 17: Performance Metrics for Bry-Boschan Business Cycle Phase: Leading Indicator Probit

	3 Months Ahead (0.3)	6 Months Ahead (0.25)	12 Months Ahead (0.35)
Accuracy	0.56	0.67	0.82
Balanced Accuracy	0.56	0.66	0.81
Sensitivity	0.6	0.8	0.78
Specificity	0.52	0.52	0.85
F1-Score	0.59	0.71	0.8

Table 18: Performance Metrics for Bry-Boschan Business Cycle Phase: Leading Indicator Random Forest

B.2.2 SARB Business Cycle Phase

	3 Months Ahead (0.2)	6 Months Ahead (0.2)	12 Months Ahead (0.35)
Accuracy	0.56	0.55	0.55
Balanced Accuracy	0.54	0.52	0.55
Sensitivity	0.72	0.72	0.56
Specificity	0.36	0.31	0.54
F1-Score	0.65	0.65	0.61

Table 19: Performance Metrics for SARB Business Cycle Phase: Leading Indicator Probit

	3 Months Ahead (0.2)	6 Months Ahead (0.2)	12 Months Ahead (0.3)
Accuracy	0.56	0.55	0.59
Balanced Accuracy	0.54	0.52	0.53
Sensitivity	0.72	0.72	0.78
Specificity	0.36	0.31	0.28
F1-Score	0.65	0.65	0.7

Table 20: Performance Metrics for SARB Business Cycle Phase: Leading Indicator Logit

	3 Months Ahead (0.3)	6 Months Ahead (0.25)	12 Months Ahead (0.3)
Accuracy	0.5	0.55	0.72
Balanced Accuracy	0.51	0.53	0.67
Sensitivity	0.44	0.61	0.89
Specificity	0.57	0.46	0.45
F1-Score	0.5	0.61	0.8

Table 21: Performance Metrics for SARB Business Cycle Phase: Leading Indicator Random Forest

B.2.3 OECD Business Cycle Phase

	3 Months Ahead (0.3)	6 Months Ahead (0.25)	12 Months Ahead (0.35)
Accuracy	0.69	0.68	0.44
Balanced Accuracy	0.70	0.67	0.44
Sensitivity	0.53	0.78	0.44
Specificity	0.88	0.56	0.44
F1-Score	0.65	0.72	0.44

Table 22: Performance Metrics for OECD Business Cycle Phase: Leading Indicator Probit

	3 Months Ahead (0.2)	6 Months Ahead (0.3)	12 Months Ahead (0.3)
Accuracy	0.68	0.65	0.47
Balanced Accuracy	0.69	0.65	0.47
Sensitivity	0.63	0.67	0.75
Specificity	0.75	0.63	0.19
F1-Score	0.69	0.67	0.59

Table 23: Performance Metrics for OECD Business Cycle Phase: Leading Indicator Logit

	3 Months Ahead (0.3)	6 Months Ahead (0.3)	12 Months Ahead (0.3)
Accuracy	0.6	0.71	0.66
Balanced Accuracy	0.61	0.71	0.66
Sensitivity	0.47	0.72	0.75
Specificity	0.75	0.68	0.56
F1-Score	0.56	0.72	0.69

Table 24: Performance Metrics for OECD Business Cycle Phase: Leading Indicator Random Forest

B.3 All Features Model Results (One-quarter ahead)

B.3.1 BBQ Business Cycle Phase

	Random Forest (0.6)	Lasso Regression (0.6)	Ridge Regression(0.55)
Accuracy	0.82	0.59	0.56
Balanced Accuracy	0.82	0.59	0.56
Sensitivity	0.85	0.45	0.55
Specificity	0.79	0.74	0.58
F1-Score	0.83	0.53	0.56

Table 25: Performance Metrics for Bry-Boschan Business Cycle Phase: All Features

B.3.2 SARB Business Cycle Phase

	Random Forest (0.5)	Lasso Regression (0.45)	Ridge Regression(0.5)
Accuracy	0.85	0.67	0.64
Balanced Accuracy	0.85	0.66	0.63
Sensitivity	0.83	0.72	0.67
Specificity	0.87	0.6	0.6
F1-Score	0.86	0.7	0.67

Table 26: Performance Metrics for SARB Cycle Phase: All Features

B.3.3 OECD Business Cycle Phase

	Random Forest (0.55)	Lasso Regression (0.6)	Ridge Regression(0.55)
Accuracy	0.83	0.75	0.72
Balanced Accuracy	0.83	0.75	0.72
Sensitivity	0.84	0.74	0.68
Specificity	0.82	0.76	0.76
F1-Score	0.84	0.76	0.72

Table 27: Performance Metrics for OECD Cycle Phase: All Features