INTERNATIONAL MONETARY FUND

IMF Country Report No. 24/289

KINGDOM OF LESOTHO

SELECTED ISSUES

September 2024

This Selected Issues paper on Kingdom of Lesotho was prepared by a staff team of the International Monetary Fund as background documentation for the periodic consultation with the member country. It is based on the information available at the time it was completed on August 22, 2024.

Copies of this report are available to the public from

International Monetary Fund • Publication Services
PO Box 92780 • Washington, D.C. 20090
Telephone: (202) 623-7430 • Fax: (202) 623-7201
E-mail: publications@imf.org Web: http://www.imf.org

International Monetary Fund
Washington, D.C.

INTERNATIONAL MONETARY FUND

KINGDOM OF LESOTHO

SELECTED ISSUES

August 22, 2024

Approved By
Andrew Tiffin (AFR)

Prepared by Adrian Alter, Samuele Centorrino, and Andrew Tiffin.

CONTENTS

APPLICATIONS OF MACHINE LEARNING IN LOW-INCOME COUNTRIES: LESOTHO[1]

In the face of the evolving global economic landscape and the inherent challenges of data scarcity in developing economies, this Selected Issues Paper (SIP) delves into few applications of machine learning (ML), with a particular application to economic forecasts in Lesotho. Amid delayed and often revised GDP data, this paper explores the potential of ML to provide real-time insights into growth and inflation trends, crucial for informed policymaking. By leveraging non-traditional data and employing a variety of ML models, the paper presents a comprehensive analysis of current economic activity, evaluates the accuracy of standard statistical measures, and forecasts future inflation trends. The findings underscore the efficacy of ML in reducing prediction errors and highlight the significant role of alternative data in circumventing the limitations posed by traditional economic indicators. This paper contributes to the broader debate on the application of advanced computational techniques in economic forecasting, offering valuable insights for policymakers in Lesotho and similar countries grappling with data constraints and the need for timely economic analysis.

A. Introduction

1. In a shock-prone world, policymakers need to adapt rapidly to steer the economy in the right direction—a task that depends on an accurate assessment of underlying trends and outlook. However, tracking economic activity in real time and informing accurate forecasts requires adequate, reliable, and timely data, which is often a challenge in developing economies. In Lesotho, quarterly GDP data is often revised significantly and released with a lag. Other high-frequency measures of economic activity such as Purchasing Managers' Indices (PMI) and employment surveys are non-existent, complicating further policymakers' job.

2. Economic activity–proxied by official GDP releases–is typically published with a delay and frequently subject to significant revisions. This has been a long-standing issue in Lesotho primarily due to the lack of coordination between various public-sector agencies and the Bureau of Statistics (BOS). Apart from capacity constraints, a further challenge for GDP data stems from an outdated base year (2012), which can give rise to added complications as the base year for consumer prices is 2022. Moreover, the overreliance on a few fiscal indicators, such as VAT receipts, to estimate activity makes it challenging to capture all aspects of the economy. Relatedly, the pandemic in 2020-21 has further affected the structure of economy, with anecdotal evidence and survey data pointing to a wider share of the informal economy. This all makes it challenging to gauge real-time economic activity.

[1] Prepared by Adrian Alter, Samuele Centorrino, and Andrew Tiffin. The paper benefited from discussions and comments received from authorities, academics, and experts in Lesotho during the 2024 Article IV consultation. Assistance with the alternative data by Hamza Mighri and Hany Abdel-Latif is gratefully acknowledged.

3. **At the same time, reliable inflation forecasts are necessary for policy purposes, including setting wage growth expectations, and guiding projections of other important variables such as fiscal forecasts.** Over the past few years, inflation has proved particularly difficult to forecast even in advanced economies, owing to a sequence of large exogenous shocks such as the Covid-19 pandemic and the war in Ukraine, which caused commodity prices to surge and disrupted global supply-chains. In such circumstances, standard time-series analysis that draws primarily on recent trends may not be a good guide, suggesting that the use of high-frequency and alternative data may help capture inflection points and structural changes more promptly.

4. **Meeting this challenge requires a flexible framework that focuses on accurate predictions.** The framework must also be able to deal with a large set of variables, potentially complex data generating processes, and limited degrees of freedom. Machine learning (ML)**,** a growing branch of applied computational statistics**,** is well-suited for the central task of nowcasting and forecasting—ML differs from traditional econometric models in its focus; rather than concentrating on inference and causality, it primarily aims at enhancing the accuracy of out-of-sample predictions.

5. **Machine-learning algorithms and applications have been a subject of interest among policymakers, gaining increasing attention over the past decade.** For instance, nowcasting models have been developed for many emerging and developing countries including Lebanon (Tiffin, 2016), sub-Saharan Africa (Barhoumi and others, 2022), India (Iyer and Gupta, 2019), and South Africa (Botha and others, 2021). Other applications include crisis prediction (IMF, 2021) and associated analysis (Tiffin, 2019). For inflation forecasting, machine learning techniques have been applied in the context of the United States (Madeiros and others, 2021), Brazil (Araujo and Gaglianone, 2023), and the euro area (Lenza and others, 2023), among others.

6. **This SIP focuses on a few applications of ML techniques that allow us to circumvent the data scarcity issues prevalent in Lesotho, by leveraging alternative datasets.** Using ML, three main questions are explored: i) What is GDP growth currently?; ii) Is economic activity under- or over- estimated by the standard statistics?; iii) Where is inflation heading? These questions are key for assessing the actual macroeconomic stance and informing policy decisions. Before delving into these applications, we'll briefly discuss the methodology.

B. Machine Learning Fundamentals

7. **There is no broadly accepted consensus on the definition of machine learning.** As a general guide, the field has its origins in computational statistics, and is chiefly concerned with the use of algorithms to identify patterns within a dataset (Kuhn and Johnson, 2018). The actual algorithms can range from the simplest OLS regression to the most-complex "deep learning" neural network; but ML is distinguished by its often single-minded focus on predictive performance—indeed, the essence of machine learning is the design of experiments to assess how well a model trained on one dataset will predict new data.

8. Fitting is easy; prediction is hard. For practitioners looking to make accurate predictions, focusing on a model's in-sample fit is insufficient. Indeed, within the machine-learning literature it is often stressed that in-sample fit generally tells us little of value, other than the number of parameters in a model (it is always possible to improve in-sample fit by adding more parameters). The key danger is that a model with a supposedly good in-sample fit for a particular dataset may in fact be modeling the idiosyncratic noise within that dataset. When taken out of sample, then, the model will perform poorly. Such a model is said to be "overfit". This is a particular danger for complex, flexible models, which is why such models have often been shunned by economists. But this is also the core issue that ML seeks to address. So, as long as the principal goal is making accurate predictions, ML makes available a much wider and richer range of estimation techniques, allowing us to more fully exploit *all* available information.

9. Estimating out-of-sample performance using in-sample data: Holdout validation. The process of predicting how a model will perform on new data is called model validation. In holdout validation, the data is split into a training and testing set. The model is built using only the training set and is then asked to make predictions using inputs from the test set. The "validation error" is the difference between these predictions and the actual test-set outcomes, and serves as a gauge of likely out-of-sample performance going forward. This error can be used to help choose between different models, or indeed, between different versions of the same type of model—in a penalized regression, for example, the weight of the penalty can be chosen in order to minimize the validation error (this process is called "tuning").

10. Estimating out-of-sample performance using in-sample data: Cross validation. As an alternative approach, cross validation takes advantage of the entire data set. The basic idea is simple: First, divide the entire dataset into K folds (say, K=3), take one of those folds and set it aside as a test set. Second, using the remaining (2) folds as a training set, estimate the model, and then use the test set to determine the model's prediction error. Finally, repeat this procedure using all combinations of the test and training sets, producing an array of (3) validation errors associated with that model, which then provides a gauge of its average out-of-sample performance. Once again, this metric can be used to help choose between different types of models. For models with pre-determined settings (for instance, the penalty weight in a penalized regression), cross validation can help determine the setting that optimizes likely out-of-sample performance. Again, in machine-learning parlance, these settings are called "hyperparameters," and are "tuned" to minimize the cross-validation error. As will be shown in this paper, cross validation as a technique allows us to compare across a wide range of models—including complex non-linear models—and to select the best version of the best model that is likely to perform well out of sample.

11. Prediction is hard; explaining a prediction is even harder. This is particularly the case for complex models, where the impact of a particular variable may depend on the value of another variable, or indeed where the final prediction is the result of an ensemble of thousands of sub models. In such circumstances there is often some tension between a model's accuracy and its interpretability. In response, various methods have been proposed to help users interpret the predictions of complex models. One simple method of determining the importance of a variable is to turn it into noise and see what happens—for instance, we can take a particular variable, scramble all of the observations to random dates and see how the fit of the model deteriorates. A large loss in explanatory power would indicate that the variable is important in making predictions. This procedure ("variable importance") is an easy-to-understand method that that provides an overall explanation of how different variables feature globally in the model.

12. **Explaining an individual prediction: Shapley Values.** Shapley values initially came out of a core question in game theory: in a team of multiple players with different skill sets, what is the fairest way to allocate a collective payoff? One solution is to imagine the players joining in sequence, and then keeping track of their marginal contribution. But what if some players, say Alice and Bob, have similar skill sets? Then, it might be the case that Alice would have a higher marginal contribution if she joined the group before Bob, as she would be the *first* one to provide their overlapping skill set. When Bob joined, his marginal contribution would be lower. The Shapley Value concept was developed in response to this problem and can be understood as finding each player's marginal contribution, averaged over every possible sequence in which the players are added to the group. To take the simplest example, suppose Alice and Bob are sharing a taxi, and Alice lives on the route to Bob's house. Their marginal contribution to the cost of the taxi ride will depend on the order in which their claims are considered. The Shapley Value, however, will average these contributions over all conceivable sequences (in this case there are only two) to arrive at the fairest possible allocation. This concept can also be used to explain the contributions of different variables to an individual prediction. Shapley Values divide this prediction (payoff) among the variables (players) in a way that fairly represents their contributions and considers all possible interactions. Since Shapley Values divide credit for joint work, the Shapley Value for a variable can be different for two dates even if that variable is unchanged, depending on what happens to other variables.

A Shapley-Value Example

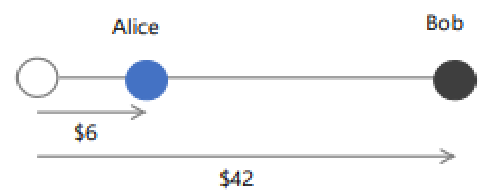

	Alice Pays	Bob Pays
Alice goes first	$6	$36
Bob goes first	$0	$42
Shapley	$\frac{\$0+\$6}{2} = \$3$	$\frac{\$36+\$42}{2} = \$39$

C. GDP Nowcasting

Predicting the Next GDP Release

13. **In this section, we employ a machine-learning approach to nowcast economic activity in Lesotho, building on the framework developed by Barhoumi and others (2022).** Again, our key question is "What is GDP growth at present?" Although the BoS does provide quarterly GDP estimates in a relatively timely fashion, these are still characterized as "experimental" and are subject to sizable revisions. So it is useful to have an independent monthly estimate of GDP trends, drawing on all available information.

14. Lesotho has a broad range of high-frequency variables that can be used to gauge GDP growth. These predictors need to be available both for a sufficiently long period of time, and released ahead of the GDP data. Equally important, they should have a historical relationship to GDP growth, either linearly or non-linearly. Note that they do not need to be *causes* of GDP growth, they just need to contain information of underlying activity (either as a cause or a symptom). Machine learning allows us to consider both traditional and alternative data, providing a list of 29 possible predictors. Traditional data includes trade-related variables (for instance, imports and exports), fiscal outturns (water royalties, government spending), and activity in South Africa (PMI, business confidence). Alternative data includes items such as Google searches on "travel to Lesotho." The correlation matrix (Figure 1) reveals a strong cluster of predictors associated with GDP growth, which includes trade-related indicators, electricity consumption, and tax revenue. A second cluster includes most of the indicators from South Africa.

Figure 1. Lesotho: Correlation Plot (Quarterly Nowcasting Model Data)

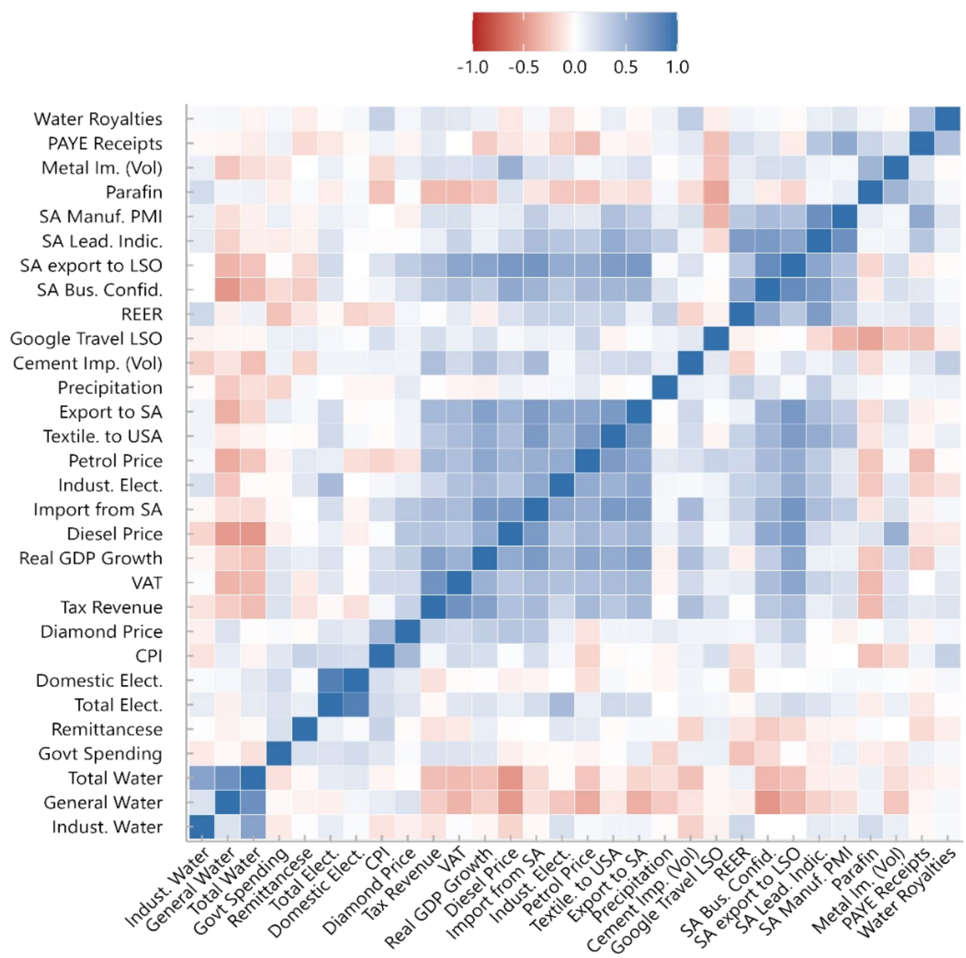

Sources: Haver Analytics, Bloomberg LLC, Google trends, National authorities, and IMF staff calculations.
Note: PAYE Receipts = Pay as you earn tax; VAT = Value added tax; REER = Real effective exchange rate; CPI = Consumer price index; SA = South Africa.

Comparing Models

15. Using the cross-validation techniques outlined above, we assess the likely out-of-sample performance of a broad range of ML models. To gauge performance, we split the data into a training set (48 quarterly observations leading up to Q4 2022) and a small testing set (4 quarterly observations in 2023). We consider eight different model types, ranging from simple OLS to a more complex neural net (Appendix A). These models are then optimized based on their validation errors in the training set, as determined from a 5-fold cross-validation framework—again, for those models with parameters that need to be chosen in advance ("hyperparameters"), these are set to minimize the average validation error. The optimal version of each model is then ranked (in a "horserace") according to its validation performance. From the eight models considered, the model most likely to do well out of sample is the Support Vector Machine (SVM) with a linear kernel (Figure 2a). As expected, given the large number of predictors, classic OLS does the worst, illustrating the advantages of a ML approach.

Figure 2. Lesotho: GDP Nowcasting Results

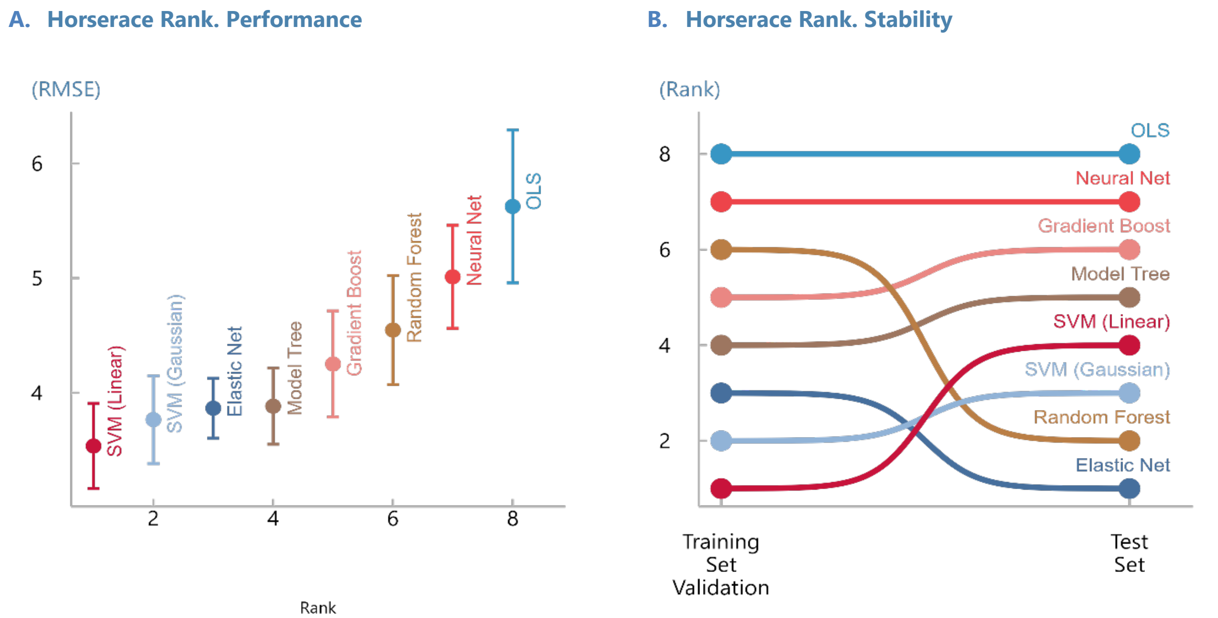

Sources: Haver Analytics, Bloomberg LLC, Google trends, National authorities, and IMF staff calculations.

Building a Stable Model: The Ensemble

16. However, the horserace rankings may not be stable. As a further experiment, we look at the ranking of our optimized models based on their performance in the small test set only. The rankings change significantly (Figure 2b), suggesting that the top-performing model may change as more data points are added. To offset this possibility, we build an ensemble model that combines the predictions of all eight models, where the weight of each model is determined optimally (again using 5-fold cross validation across the training set). This weighted average limits the noise from any individual model and so adds to the stability of our nowcasts—the optimal ensemble turns out to be a weighted average of five types of model, with most of the weight given to the SVM (linear) model and the Extreme Gradient Boosting ("XGBoost") model.[2] The nowcast predictions from the ensemble model are shown in Figure 3. The figure also provides variable-importance measures for the most important high-frequency variables. In this case, the leading indicator of activity in South Africa seems to be most important predictor of GDP growth in Lesotho (in that scrambling this variable does the most damage to the model's performance). Other important variables include imports from South Africa, tax revenues, and textile exports to the USA.

Figure 3. Lesotho: Ensemble Nowcasting Model

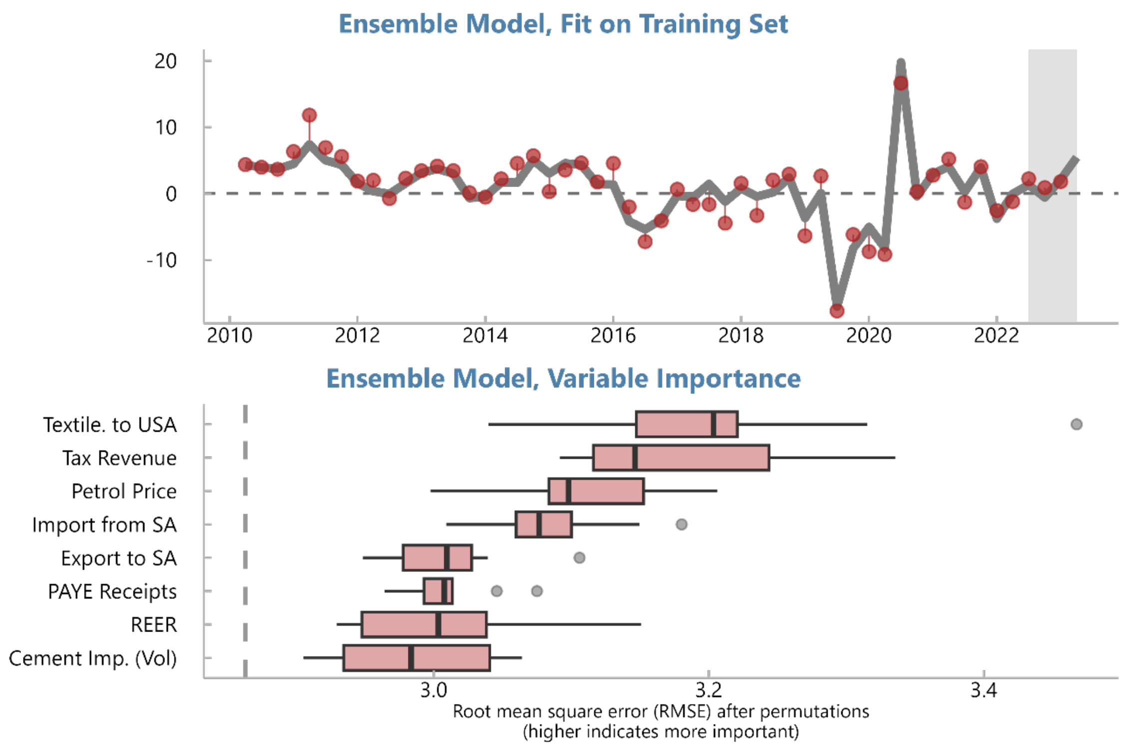

Sources: Haver Analytics, Bloomberg LLC, Google trends, National authorities, and IMF staff calculations.

Results

17. The final model suggests that Lesotho's real GDP growth improved significantly in Q1-2024, after bottoming out in Q3-2023. The nowcast shows an uptick in annual growth from 0.5 percent at end-2023 to 1.8 percent in March 2024 (Figure 4). The particular drivers of the latest monthly prediction (March 2024) can be seen in the Shapley Value breakdown of that month's nowcast. Looking at the factors that pushed the prediction above the sample average, the main upward impetus was associated with a strong boost in water royalties, and an improvement in textile

[2] The combined ensemble is constructed using a Lasso regression of GDP growth, in which the predictions from a pool of candidate models are used as variables in the regression. The pool of candidates is broader than the eight best versions of the different model types, and instead includes a range of top-performing hyperparameter combinations for each model. For a more detailed discussion of this ensemble methodology, see Kuhn and Silge (2022).

exports to the USA. On the other hand, lower exports to South Africa helped bring the prediction back down.

Figure 4. Lesotho: Monthly Nowcast of GDP Growth (2015–2024)

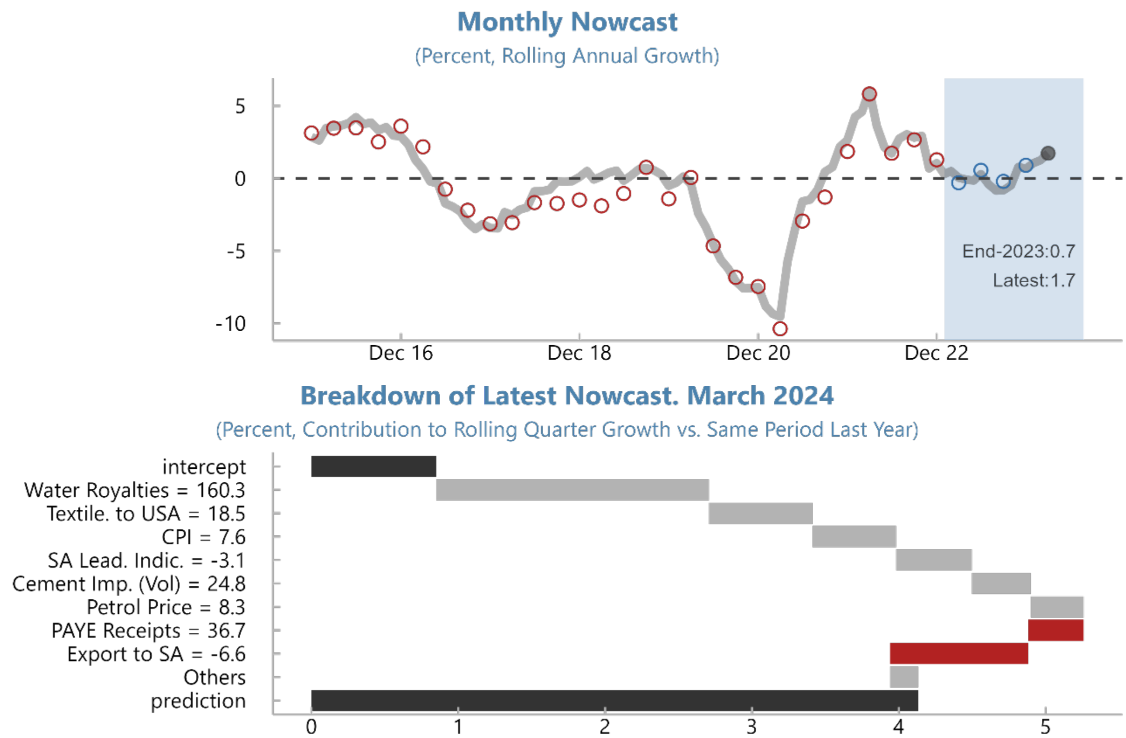

Sources: Haver Analytics, Bloomberg LLC, Google trends, National authorities, and IMF staff calculations.

D. GDP Evaluation

18. In this section, real GDP series are empirically evaluated using historical quarterly data. The primary motivation of this exercise is the potential misestimation of real economic activity by Lesotho's Bureau of Statistics due to lack of data, capacity constraints, and other barriers common in developing countries. Second, the Covid-19 pandemic was a significant shock which may have notably impacted the informal economy, thus potentially changing previous sensitivities (e.g., to taxes, consumption, etc.).

19. Combining a comprehensive set of variables with ML techniques provides alternative GDP estimates. The selected ML technique is the LASSO regression, allowing to establish relationships with selected regressors in the context of limited quarterly data. The relationships are estimated using 40 observations covering the period prior to the pandemic (i.e., Q1-2010 – Q4-2019), with a goal of tracking closely the actual series (Figure 5.A.). This is validated by the in-sample fit which is close to 1, with an R^2 higher than 0.9.

20. The set of regressors covers standard macroeconomic indicators, fiscal variables, and alternative series such as night lights, air pollution, and precipitation. The most relevant and significant variables selected by LASSO are related to fiscal outturns, validating the large footprint of the government in the economy. For instance, government's goods and services variable has the highest standardized coefficient (2.4), followed by water royalties and fiscal balance. In addition, water consumption and imports of construction materials are equally relevant. From the alternative sources, air pollution and Google searches were selected by the LASSO regression, avoiding overfitting and multicollinearity.

21. The results point to a potential underestimation of real GDP close to 9 percent at end-2023. Using the relationships suggested by the LASSO regression, the predicted "out-of-sample" GDP series deviate significantly from the actual data over the post-pandemic period (e.g., 2020-2023). The largest deviation is found in 2022, with the estimated series being about 14 percent greater than the actual, potentially reflecting the lack of post-pandemic rebound and issues stemming from a larger informal sector (Figure 5.B).

Figure 5. Lesotho: Real GDP Evaluation

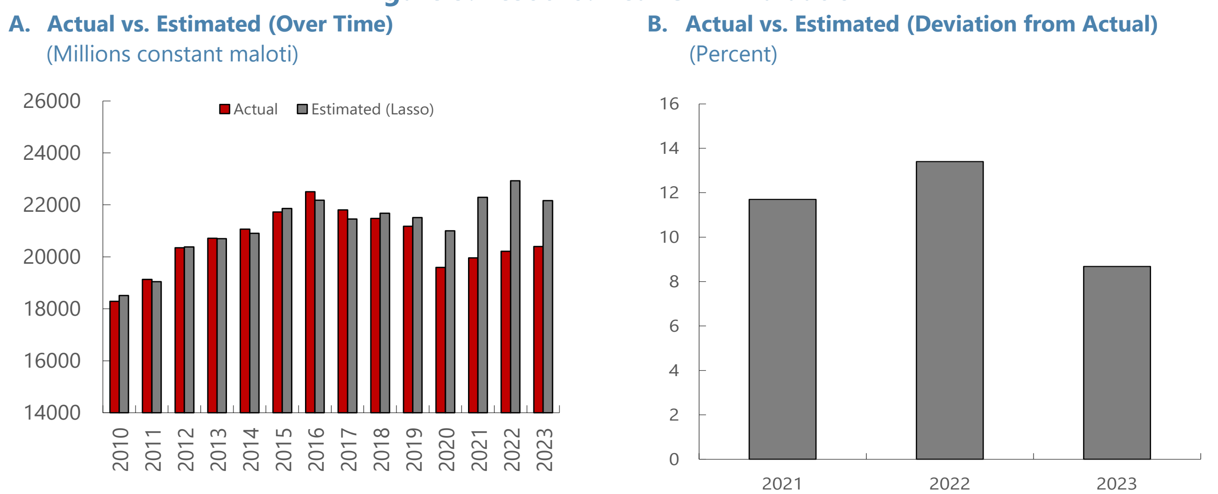

Sources: Haver Analytics, Bloomberg LLC, Google trends, National authorities, and IMF staff calculations.

E. Inflation Forecasting

22. In this section, ML techniques are applied to forecasting inflation in Lesotho and South Africa, leveraging on the rich information available at the global, regional, and local level. The motivation for improving inflation forecasts stems from several sources. First, inflation has surprised to the upside in recent years, due to multiple and sequential exogenous shocks such as demand resurgence in the aftermath of the pandemic, higher food/cereal prices because of the war in Ukraine, and supply-chain disruptions on the back of rising geopolitical fragmentation. Second, inflation rates in Lesotho and South Africa have significantly diverged in 2023–24 primarily due to a larger share of food items in Lesotho's CPI basket. While inflation trends were widening in 2023, they started to slowly converge in 2024. The pace of this convergence is ultimately important for monetary policy decisions. Finally, systematic forecasting errors may result in a misestimation of fiscal and other variables, potentially leading to welfare costs.

23. For this exercise, we focus on two flagship models, each with different strengths and weaknesses.

- *Random Forests* is a widely used example of an ensemble technique, where the focus is on improving predictive performance by combining the predictions of a (very) large number of separate models, and so eliminating idiosyncratic noise through the law of large numbers. Random forest models are often superior in capturing potential sources of nonlinearity.

- *Elastic Net Regression*, on the other hand, is a common example of a regularization technique, where the penalty term in the regression reduces the model's complexity and so prevents overfitting directly. Elastic net models are general and parsimonious, leading to both easily explainable results, and a model that also performs well in flexibly detecting turning points.

24. For our purposes, regularization techniques are used in a standard linear model without interaction terms or other nonlinearities, supporting simplicity and easy-to-explain results. At the same time, ensemble methods are used to provide a flavor of the importance of nonlinearities in forecasting inflation (see, e.g., Lenza and others, 2023). Both models are well positioned to deal with a large set of potential predictors while solving the challenge of limited degrees-of-freedom, with the series for inflation in Lesotho having only 70 quarterly observations. Four lags are included for each predictor to allow for delayed passthrough, along with inflation lags.

25. In both cases, the goal is to find the best model for forecasting inflation four quarters ahead. Thus, the CV scheme is constructed accordingly. We start training the model on the first five years of data and evaluate it in the next year. At each subsequent evaluation, we expand the training set by one year and shift the testing set by one year, until we reach the end of the sample. As the sample size is equal to 70 quarters, this allows for 12 different training and testing sets. The tuning parameters are chosen to minimize the Mean Absolute Scaled Error (MASE), a well-accepted measure of forecasting accuracy (Hyndman and Koehler, 2006).

26. The large set of regressors includes global, regional, and local variables. Although this set of potential predictors is not exhaustive, the main global indices include commodity prices such as oil and international food prices, as well as the Baltic Freight Index–a proxy for worldwide transportation costs by sea. International prices for food and oil are obtained from the WEO database, which additionally offers forecasts for the upcoming quarters, based on the futures market. At the regional level, the rich information from South Africa is leveraged to predict inflation in both countries. For instance, South Africa's inflation expectations one- and two-years ahead are utilized. Based on existing literature, the exchange rate typically plays an important role in driving inflation, often with a delay, and thus both nominal exchange rates and Rand to USD are included. To control for economic activity, and implicitly for demand factors, the predictor set includes business satisfaction, a composite leading indicator, PMIs, and motor vehicle sales, all sourced from Haver Analytics.

27. The out-of-sample performance finds the elastic net model to have the lowest RMSE compared to other models. The Random Forest and Elastic Net models are benchmarked against a standard autoregressive model with 4 lags (i.e., AR(4)) for both South Africa and Lesotho (Figure 6). The results for Lesotho point to a substantial improvement in the RMSE compared to the benchmark model, with a reduction of about 30 percent at the median for the Elastic Net. The variables selected by the Elastic net model include South Africa's inflation expectations, the leading indicator, and the PMI, along with international food prices (Table 2). In contrast, the out-of-sample performance improves slightly less in South Africa, with the median RMSE just under 1 for the Elastic Net model. The primary variables selected in this case are akin to those for Lesotho, barring the inclusion of business satisfaction and global oil prices.

28. The inflation forecasts indicate relatively stable inflation for South Africa and a moderate decline in Lesotho, signaling further convergence in rates over the next four quarters. In South Africa, the one-year ahead predictions vary between 4.9 and 6.2 percent across the three models (Figure 7.B). However, the prediction of the model with lowest RMSE (i.e., the Elastic Net) is close to 5 percent at the end of Q1-2025. At the same time, the one-year ahead forecasts for Lesotho range from 5.6 to 6.4 percent, with the forecast from the Elastic Net model closely tracking the Random Forest and ending FY24/25 at 5.9 percent (Figure 7.A). Importantly, these models consistently predict a narrower divergence between the inflation trends in South Africa and Lesotho.

Figure 6. Lesotho: Inflation Forecasting: Out-of-Sample Performance

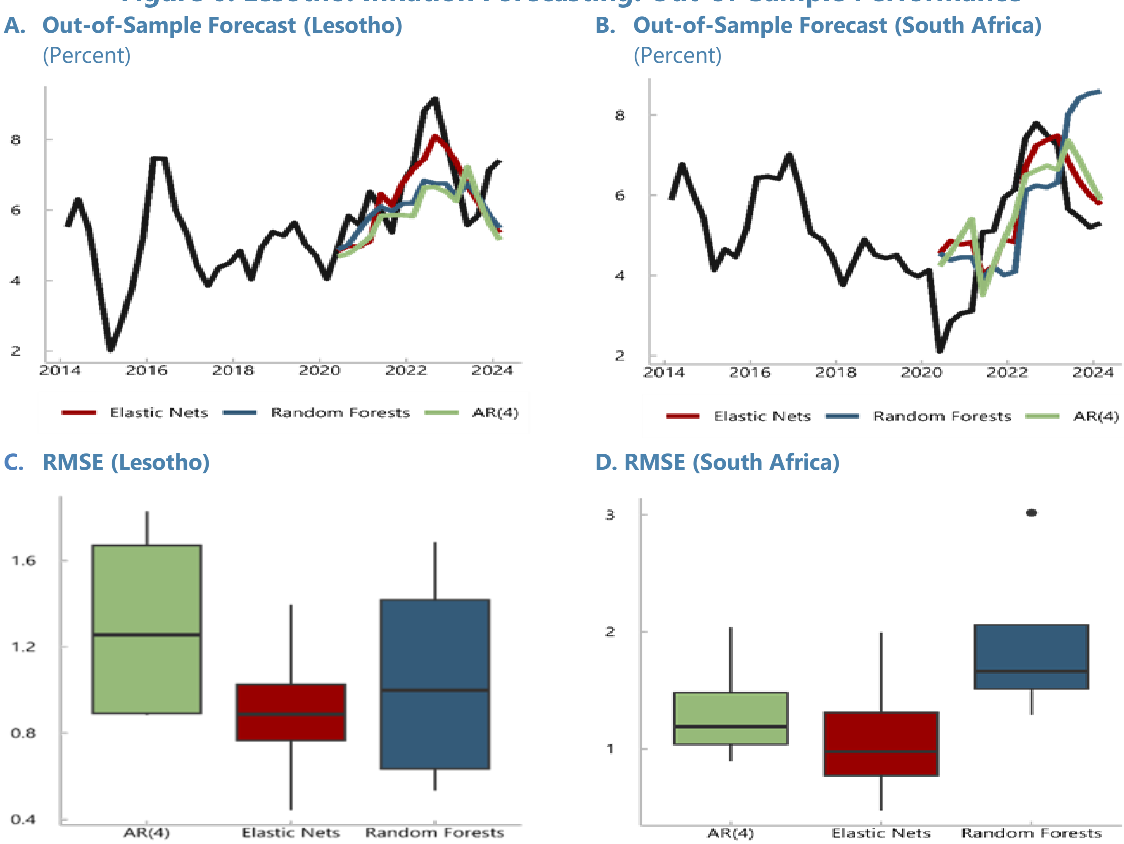

Sources: Haver Analytics, Bloomberg LLC, Google trends, National authorities, and IMF staff calculations.
Note: RMSE = Root Mean Squared Errors; AR(4) = autoregressive model with 4 lags.

Figure 7. Lesotho: Inflation Forecasts (One-year Ahead)

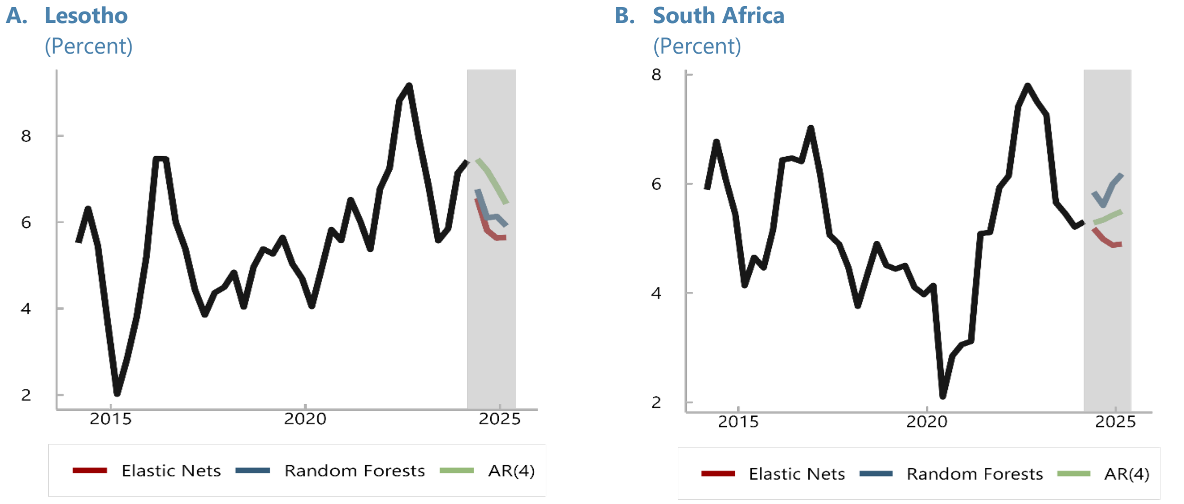

Sources: Haver Analytics, Bloomberg LLC, Google trends, National authorities, and IMF staff calculations.

F. Conclusion

29. The ability of policymakers to adapt and guide the economy hinges on accurate and timely data, a challenge particularly pronounced in developing economies like Lesotho. Policymakers face difficulties in real-time tracking of economic activity due to delayed and often revised quarterly GDP data, compounded by the absence of high-frequency economic activity measures. This situation is exacerbated by coordination issues among public-sector agencies and an outdated GDP base, making it challenging to capture the full spectrum of economic activity, especially in the wake of the pandemic which has likely expanded the informal economy's share.

30. Machine learning techniques emerge as a potent solution to these challenges, offering a flexible framework for nowcasting and forecasting by focusing on prediction accuracy over causality. These techniques, including LASSO regression, Elastic Net, Random Forests, and Support Vector Regressions, have been applied in various countries to enhance forecasts of economic activity and inflation. In Lesotho, ML helps circumvent data scarcity, leveraging alternative datasets to provide real-time insights into GDP growth, the accuracy of standard statistics, and future inflation trends, thus aiding in informed policymaking.

31. The application of ML in Lesotho reveals promising results, with ensemble models predicting an uptick in real GDP growth and providing more accurate inflation forecasts. This approach not only aids in navigating immediate economic challenges but also in evaluating the potential bias in GDP estimation and improving inflation forecasts amidst global and regional shocks. By embracing ML and alternative data, Lesotho can achieve more reliable economic insights, guiding policy decisions towards sustainable growth and development.

References

Araujo, G.S. and Gaglianone, W.P. 2023. "Machine learning methods for inflation forecasting in Brazil: New contenders versus classical models." *Latin American Journal of Central Banking*, *4*(2), p.100087.

Barhoumi, K., Choi, S.M., Iyer, T., Li, J., Ouattara, F., Tiffin, A., and Yao, J. 2022. "Overcoming data sparsity: A Machine Learning approach to track the real-time impact of COVID-19 in Sub-Saharan Africa." *International Monetary Fund*, WP/22/88.

Botha, B., Reid, G., Olds, T., Steenkamp, D. and van Jaarsveld, R. 2021. "Nowcasting South African GDP using a suite of statistical models." Economic Research and Statistics Department, South African Reserve Bank.

Breiman, L. 2001. "Random forests." *Machine Learning Journal*, 45, pp. 5–32.

Green, P.J., and Silverman, B.W. 1993. "Nonparametric Regression and Generalized Linear Models: A roughness penalty approach." *Taylor & Francis*.

Hyndman, R. J. and Koehler A. B. 2006. "Another look at measures of forecast accuracy." *International Journal of Forecasting*, 22(4), pp. 679–688.

Iyer, T. and Gupta, A.S. 2019. "Nowcasting economic growth in India: The role of rainfall." *Asian Development Bank Economics WP Series*, (593).

IMF. 2021. "How to Assess Country Risk: The Vulnerability Exercise Approach Using Machine Learning," *Technical Notes and Manuals (International Monetary Fund)*, TNM/21/03.

Kuhn, M., and Johnson, K. 2018. "Applied Predictive Modeling," 1st ed. 2013, Corr. 2nd printing 2018, *Springer Texts in Statistics*.

Kuhn, M., and Silge, J. 2022. "Tidy Modeling with R," *O'Reilly Publishing*.

James, G., Witten, D., Hastie, T., Tibshirani, R. 2023. "An Introduction to Statistical Learning: with Applications in R." 2nd Ed., *Springer Texts in Statistics*.

Lenza, M., Moutachaker, I. and Paredes, J. 2023. "Density Forecasts of Inflation: A Quantile Regression Forest Approach", Working Paper Series, No 2830, ECB.

Lundberg, S.M., and Lee, S. 2017. "A unified approach to interpreting model predictions." *Advances in Neural Information Processing Systems.*

Medeiros, M.C., Vasconcelos, G.F., Veiga, Á. and Zilberman, E. 2021. Forecasting inflation in a data-rich environment: the benefits of machine learning methods. *Journal of Business & Economic Statistics*, *39*(1), pp.98–119.

Ribeiro, M.T., Singh, S., and Guestrin, C. 2016. "Why should I trust you?: Explaining the predictions of any classifier." *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. ACM.*

Tibshirani, R. 1996. "Regression Shrinkage and Selection via the Lasso." *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1), 267–288.

Tiffin, A. 2016. "Seeing in the dark: A machine-learning approach to nowcasting in Lebanon," *International Monetary Fund*, WP/16/56.

Tiffin, A. 2019. "Machine Learning and Causality: The Impact of Financial Crises on Growth," *International Monetary Fund*, WP/258.

Zou, H., Hastie, T. 2005. "Regularization and Variable Selection Via the Elastic Net." *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67(2), pp. 301–320.

Appendix I. A Brief Summary of Different Machine Learning Models

1. **Ordinary Least Squares** (OLS): OLS is the classic linear regression method. It fits a straight line through data points by minimizing the sum of squared residuals. OLS assumes that errors are normally distributed and homoscedastic. It's simple, interpretable, and widely used in econometrics, but does not always do well with large numbers of predictors.

2. **Elastic Net:** Elastic Net adds a penalty term to the OLS regression. It combines the Lasso penalty (which selects the best predictors) and the Ridge penalty (which nudges the coefficient of each predictor towards zero), allowing it to select relevant features while also shrinking coefficients. As a standard workhorse for dealing with a large number of predictors, it's a linear regression variant that nicely balances feature selection and handles multicollinearity well.

3. **Support Vector Machine** (SVM): Unlike OLS, the SVM (regression) uses the ε-insensitive loss function instead of the standard squared loss. The best-fit is represented by a hyperplane that maximizes the number of data points within the ε-tube around it. Essentially, this tube is a "margin of tolerance." The model is efficient, includes regularization to reduce overfitting, and works well even in high-dimensional spaces. Potential non-linearities are addressed by projecting the data into a higher-dimensional representation through a pre-specified kernel function ("the kernel trick"). Kernel variants considered in this paper include: the linear kernel, and the Gaussian kernel ("Radial Basis Function").

4. **Random Forest:** The Random Forest algorithm builds an ensemble of side-by-side decision trees (the "forest") and aggregates their predictions. Each tree is trained on a random subset of features and data points, reducing overfitting. It's robust, interpretable, and handles missing values well.

5. **Extreme Gradient Boosting** (XGBoost): Unlike Random Forests, XGBoost builds its ensemble of decision trees sequentially, with each new tree correcting the cumulative errors made by all the previous trees ("boosting"). XGBoost is widely used in competitions and real-world applications.

6. **Model Tree** ("Cubist"): The Model Tree algorithm combines decision trees with linear models. Each final leaf of the tree corresponds to a different (simple) linear regression, making it highly interpretable. It's useful when relationships are nonlinear but can be approximated locally.

7. **Neural Network** (MLP): This version used in this paper is a multilevel perceptron (MLP), which is a straightforward feedforward neural network. It stacks layers of interconnected neurons, where each neuron is fed into a nonlinear activation function. These allow the network to learn complex patterns in the data, transforming a linear computations into a potentially rich, expressive representation.

8. A more technical treatment of these algorithms is beyond the scope of this paper, but details are available in James and others (2023), and Kuhn and Johnson, (2018). An accessible hands-on guide to implementing these models in practice can be found in Kuhn and Silge (2022).

Table 1. Lesotho: Real GDP Evaluation – Selected Variables

Variable	Selected	Variable	Selected
Enhanced Vegetation Index		Electricity Consumption (Total)	✓
Night Lights		Electricity Consumption (Industrial)	✓
NO2 Emissions	✓	Electricity Consumption (Households)	
Fiscal Balance	✓	Water Consumption (Total)	
Government Revenues (Total)		Exports of Textiles (to the US)	
Tax Revenues		Cement Imports	✓
SACU Transfers		Imports from South Africa	✓
VAT receits		Metal Prod. Imports	✓
Water Royalties (Total)	✓	Payemillions	
Government Expenses		Petrol Price	
Employees Compensation (gov.)		Diesel Price	
Goods and Services (gov.)	✓	Parafine	✓
Nonfinancial Assets		Total exports	✓
Google Travel Trends (LSO)	✓	Trend	
Precipications		Fixed Effects (quarter)	✓

Sources: Central Bank of Lesotho; Haver Analytics, Lesotho's Bureau of Statistics; WEO database; IMF staff calculations.

Table 2. Lesotho: Inflation Forecasting – Selected Variables

	Lesotho	South Africa
Inflation (lag 1)	✓	✓
Inflation Expectations 2-yr ahead (lag 5)	✓	
Inflation Expectations 2-yr ahead (lag 6)	✓	
Inflation Expectations 1-yr ahead (lag 7)		✓
International Food Prices (lag 2)		✓
International Food Prices (lag 3)	✓	
International Food Prices (lag 4)	✓	
International Oil Prices (lag 1)		✓
ZAF Leading Indicator (lag 5)	✓	
ZAF Leading Indicator (lag 6)	✓	✓
ZAF Business Safisfaction (lag 5)	✓	✓
ZAF Business Safisfaction (lag 7)		✓
ZAF PMI (lag 5)		✓
ZAF PMI (lag 6)	✓	

Sources: Lesotho's Bureau of Statistics; Central Bank of Lesotho; Haver Analytics, WEO database; IMF staff calculations.
Note: ZAF = South Africa; PMI = Purchasing Managers Index.