INTERNATIONAL MONETARY FUND

Nowcasting Growth Using the Bayesian Structural Time Series Model: Application to Tanzania

Prepared by Sunwoo Lee

WP/26/49

2026

MAR

WORKING PAPER

 WP/26/49

IMF Working Paper
African Department

Nowcasting Growth Using the Bayesian Structural Time Series Model: Application to Tanzania
Prepared by Sunwoo Lee

Authorized for distribution by Justin Tyson
March 2026

ABSTRACT: In light of recent global shocks and rising external volatility, there is a growing need to effectively monitor short-term economic fluctuations, especially in countries with limited access to high-frequency growth data. This paper examines the application of the Bayesian Structural Time Series (BSTS) model to the case of nowcasting quarterly economic growth in Tanzania, leveraging a range of high-frequency economic indicators. The BSTS model provides a flexible framework that incorporates trends, seasonal variations, and regression effects, while its spike-and-slab variable selection helps identify relevant indicators. This paper outlines a framework for model selection and evaluation, including robustness checks and sensitivity analysis, and demonstrate the model's relative performance. Additionally, the model's capacity to adapt to longer forecast horizons and dynamic regressors enhances its utility for understanding growth trends in changing economic environments.

RECOMMENDED CITATION: Lee, Sunwoo, 2026, "Nowcasting Growth Using the Bayesian Structural Time Series Model: Application to Tanzania". IMF Working Paper WP/26/49.

JEL Classification Numbers:	C53, E37, O11
Keywords:	Nowcasting; Bayesian models; economic activity, GDP; low-income countries
Authors' email addresses:	slee8@imf.org

WORKING PAPERS

Nowcasting Growth Using the Bayesian Structural Time Series Model: Application to Tanzania

Prepared by Sunwoo Lee [1]

[1] The author would like to thank Haris Tsangarides for his guidance; Mika Saito, Mercedes Van Martin, and the AFR Research Therapy seminar participants for their helpful comments. All remaining errors and omissions are the author's own.

Contents

I. Introduction

Amid recent global shocks and increased external volatility affecting economies, there is a growing need for monitoring short-term economic fluctuations. However, high-frequency growth data series are not readily available in many countries. Official quarterly Gross Domestic Product (GDP) figures are absent in over 60 countries[2] and are often subject to delays in others. In sub-Saharan Africa, in particular, the lag exceeds one quarter for about 20 countries.[3] These extended delays can pose challenges for timely assessments of the economy.

At the same time, an increasing number of high-frequency economic indicators have become available. New technologies have introduced novel forms of data, expanding the pool of relevant information sources. Furthermore, more countries have begun to collect and publish various statistics relevant to economic activities, ranging from the production and consumption of manufacturing inputs to the number of tourist arrivals, in addition to traditional economic indicators. Improved data availability can help gauge the status of economic activities at a higher frequency than that of official data releases.

Nowcasting economic growth using high-frequency indicators has thus become a popular option to complement official statistics that are typically published at lower frequencies or with delays. Among these tools, univariate models, such as the Bridge and Mixed Data Sampling (MIDAS) models, are favored for their simplicity and effectiveness in capturing short-term fluctuations with a limited set of indicators (Clements and Galvao (2007); Armesto et al. (2010)). Multivariate approaches, including Vector Auto-regressive Models (VAR) and Dynamic Factor Models (DFM), offer greater complexity by accounting for interdependencies among multiple economic variables, making them suitable for analyzing larger datasets (D'Agostino et al. (2010); Mariano and Murasawa (2010); Giannone et al. (2008); Banbura et al. (2011)). DFMs are particularly useful for handling high-dimensional data and addressing non-fundamental shocks. Recent advances in machine learning techniques have enabled nowcasting tools to model nonlinear relationships and efficiently process big data, thereby enhancing the accuracy and reliability of economic forecasts (Barhoumi et al. (2022); Woloszko (2020)).

Introduced by Scott and Varian (2014), the Bayesian Structural Time Series (BSTS) model provides an additional tool for time series nowcasting and forecasting. Its modular state-space framework allows for the flexible incorporation of trend, seasonality, and regression components. The use of spike-and-slab priors for variable selection enables automatic variable selection, highlighting the most relevant indicators without relying on latent factor estimation.

In developing economies – where data availability may be limited, growth dynamics are evolving, and clear policy communication is essential – the model can offer several advantages. It can accommodate sparse and irregular growth data, and its modularity provides a transparent narrative of what drives the nowcast. As new data accumulates, the model gradually shifts weight toward those that become more informative over time. The availability of full posterior distributions, rather than point estimates, allows for richer uncertainty analysis and more informed decision-making.

In contexts where formal nowcasting frameworks are scarce and prior efforts have relied on extrapolation or simple autoregressive models, this approach can thus be a flexible alternative. It can be further adapted to longer forecast horizons and to changing economic relationships through the inclusion of dynamic regressors,

[2] See Bear, Guerreiro, and Silungwe (2022).

[3] See Akbal, Choi, Narita, and Yao (2023).

which allow predictor–outcome relationships to evolve over time. Moreover, the model provides direct attribution to observable indicators, making it easier to explain drivers of the forecast output and to detect potential shifts in sectoral relevance.

This paper presents a practical framework for model selection and specification using the BSTS tool, and assesses its performance in the context of data sparsity and limited visibility into the drivers of economic activity. Building on previous applications of the BSTS model for economic forecasting (Pérez (2018); Kohns and Bhattacharjee (2023)), as well as existing initiatives to nowcast Tanzania's growth as part of broader Sub-Saharan Africa (Barhoumi et al. (2022); Akbal et al. (2023)) and by the central bank under its forecasting framework, the paper applies the approach to estimate quarterly GDP growth in Tanzania. The model's flexibility, while adding to its adaptability, necessitates careful specification to ensure reliable performance. Bayesian model averaging during the estimation process as well as averaging of selected models in the framework improve robustness.

The rest of this paper is organized as follows: Section 2 provides an overview of the model components; Section 3 outlines the methodology and application to Tanzania; Section 4 describes the model outputs; Section 5 illustrates the extensions with multi-period forecasts and dynamic regressors; and Section 6 concludes.

II. Model Components and Features

The BSTS model is a state-space model consisting of two sets of equations.[4] The observational equation sets up the relationship between the response variable and the predictors and latent variables. State equations govern the evolution of model parameters over time and are determined based on trend and seasonality assumptions. Eq (1) illustrates an example of the state space representation of a simple model, assuming random walk observed in noise:

$$\begin{aligned} y_t &= \mu_t + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \\ \mu_t &= \mu_{t-1} + u_t, \qquad u_t \sim N(0, \sigma_u^2) \end{aligned} \tag{1}$$

Any prior information or beliefs provided as model assumptions shape the prior distribution of the model parameters. Trend assumptions determine the prior distribution of parameters relevant to the trend equation which can take various forms depending on the specific features of the data and the patterns observed in the time series. Two examples of trend assumptions are presented in Section 3. The model also allows for assumptions on other components, including seasonality and regressors to capture the underlying structure of the data accurately.

The model has two key features. First, the model uses the Kalman filter for estimation. Once the user enters prior beliefs as model inputs, the algorithm derives the fitted value of the response variable using the prior distribution and subsequently updates the parameter distribution based on the errors. The simulation step is iterated to obtain the posterior distribution, which is then used to generate predictions in forecast periods.

[4] See Scott and Varian (2014) and Scott and Varian (2015) for more details.

The use of Kalman filter adds to the flexibility of the model and helps in addressing irregularity of the response variable.

Second, the model selects its regressors using the Spike and Slab method. The method is based on the prior belief that a majority of regression coefficients are precisely zero, therefore assigning a positive probability to the coefficient being exactly zero. The inclusion probability of each regressor is updated upon observing the fitted data. The expected number of predictors with nonzero coefficient can be specified to determine the prior distribution of spike and slab parameters. For instance, if the expected model size is set at 1 without any additional specific prior assumption on coefficients, the prior inclusion probability of n indicators would be set at 1/n for each. The variable selection method, coupled with Bayesian model averaging, helps mitigate issues with spurious relationships.

As outputs, the model produces a distribution of estimates rather than a single point estimate, which allows users to utilize the entire distribution. Based on their specific needs, the users can choose to focus on summary statistics such as the mean, median, or a range of estimates, enabling them to make more informed decisions regarding economic forecasts.

The model therefore is well-suited to settings where the indicator set includes noisy or weakly informative series and where transparency in model behavior is needed. Unlike factor-based approaches, which summarize large datasets through a small number of latent factors, BSTS relies on spike-and-slab priors to identify regressors that carry short-term informational content. This allows the model to filter out weakly informative or noisy series, while retaining attribution to individual indicators that influence near-term movements in the nowcast. The structural decomposition into trend, seasonal, and regression components, together with full predictive distributions, provides a transparent narrative of the drivers of the nowcast and a richer uncertainty assessment. Results from the Tanzania application in Section IV suggest that well-calibrated model specifications yield robust forecasting performance.

These advantages come with trade-offs. Computational demands increase as the number of predictors grows, in contrast to DFMs, which scale efficiently to high-dimensional environments and can perform better when the underlying data-generating process is dense. DFMs also remain more suitable in contexts with substantial data irregularities. The spike-and-slab prior used for variable selection assumes sparsity, which may lead to the exclusion of potentially informative indicators, especially in short sample periods. This can affect the model’s ability to fully capture sectoral dynamics or emerging trends. In addition, BSTS forecasts can be sensitive to prior assumptions and therefore require calibration informed by economic judgment and familiarity with the data.

Taken together, BSTS complements the other nowcasting approaches. It is most effective when interpretability, indicator-level attribution, and transparency are central to the nowcasting exercise, and when the indicator set is of moderate size so that sparsity-based selection can operate effectively. In contrast, alternative methods such as DFMs may remain preferable for very large datasets driven by broad co-movements across indicators.

III. Model Inputs and Methodology

This section presents the data inputs and methodological framework used in the nowcasting exercise. It details the selection and preprocessing of indicators and the implementation of the BSTS model in the Tanzanian context.

Model Inputs and Data Description

The nowcasting exercise requires several assumptions on model parameters in addition to a set of economic indicators as model inputs.

Trend and seasonality assumptions: A local linear trend assumption, along with a quarterly seasonality assumption, was tested for the nowcasting exercise.[5]

The linear local trend assumes that the mean and slope of the trend follow random walks. Additional assumptions can be applied to the standard deviation parameters σ_u and σ_v as well. The state space representation with the linear local trend assumption is as follows:

$$\begin{gathered} y_t = \mu_t + \tau_t + \beta' x_t + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \\ \mu_t = \mu_{t-1} + \delta_{t-1} + u_t, \qquad u_t \sim N(0, \sigma_u^2) \\ \delta_t = \delta_{t-1} + \nu_t, \qquad \nu_t \sim N(0, \sigma_\nu^2) \\ \tau_t = -\sum_{s=1}^{S-1} \tau_{t-s} + \omega_t \end{gathered} \tag{2}$$

where y_t is the variable of interest, μ_t is the level of the trend component, δ_t is the slope of the trend component, τ_t is the seasonal component ($S = 4$ for quarterly seasonality assumption), and x_t is the vector of indicators.

When applicable, the robust local linear trend (also known as the student local linear trend) provides greater flexibility in the tail assumptions. The trend assumption characterized by the below equations:

$$\begin{gathered} \mu_t = \mu_{t-1} + \delta_{t-1} + \epsilon_t, \qquad \epsilon_t \sim T_{\nu_\mu}(0, \sigma_\mu^2) \\ \delta_t = \delta_{t-1} + \eta_t, \qquad \eta_t \sim T_{\nu_\delta}(0, \sigma_\delta^2) \end{gathered} \tag{3}$$

where ν_μ and ν_δ are the tail thickness parameters of the T distribution for errors, thereby offering greater resilience to outliers. Assumptions can be made to the prior distributions by specifying the standard deviation parameters and the tail thickness parameters.

Economic Indicators and Expected model size (k): High-frequency indicators relevant to the response variable can be added to the model as regressors. The model input *expected model size (k)* determines the prior inclusion probability of the indicators. With no additional assumption made, $0 < k < 1$ is the prior

[5] Another readily available trend assumption in the BSTS R package (Scott (2024)) is the semi-local linear trend, which incorporates a change in the slope component as follows: $\delta_t = D + \phi * (\delta_{(t-1)} - D) + \nu_t, \nu_t \sim N(0, \sigma_\nu^2)$. The semi-local linear trend adds a long run slope parameter D and can be useful for long-term forecasts.

probability of each regression coefficient being non-zero. Alternative prior assumptions on inclusion probabilities can also be specified if it is preferred to include one or more key indicators as regressors with certain probabilities. Setting an informative prior can also help address overfitting issues that may result from including too many regressors in a short sample period.

To illustrate a simple case of a nowcasting model that includes multiple regressors, Figure 1 presents the average contribution of each model component to the one-period-ahead predictions of growth in Tanzania's agricultural activities. The model specification includes a local linear trend component, a quarterly seasonality component, and 130 series of regional level Normalized Difference Vegetation Index (NDVI)6 as regressors. When the expected number of regressors with non-zero coefficients is smaller, relatively more weight is put on the trend and seasonality components of the model. In the illustrated case, the trend and seasonality component move in similar way with both values of k over the sample period, while the magnitude of contribution from regressors is much smaller in the case of a lower value of k. In addition to serving as a model input that reflects the prior information regarding the informativeness of the regressors, the expected model size can also determine the model's flexibility, given the assumptions related to other model inputs. For instance, if the trend assumptions allow for limited volatility, the inclusion of additional regressors can capture unexpected fluctuation in growth that is not fully reflected by the trend component. Conversely, if the trend component allows for larger fluctuations, the additional input from regressors can help moderate the extent of volatility based on the movement of relevant regressors.

Figure 1. Contribution to predictions by components, different priors of expected model size *k*

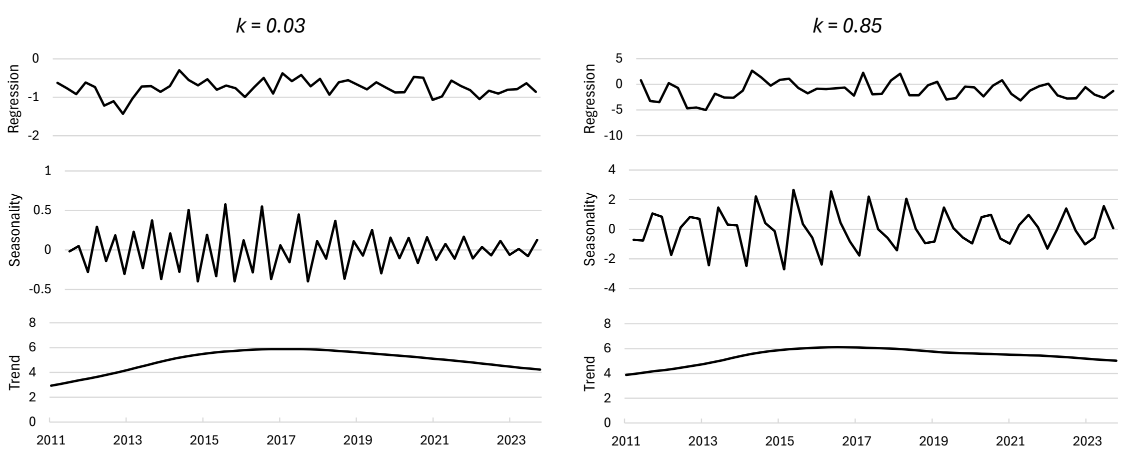

Note: 0<*k*<*1* is the expected share of regressors with nonzero coefficients.

Model Selection

A set of model specifications, based on inputs, is tested using out-of-sample Root Mean Squared Errors (RMSEs) to assess predictive performance and generate nowcast estimates. In addition to the range of expected model sizes, the default grid of parameters establishes the range of priors for standard deviation

[6] Normalized Difference Vegetation Index is a measure used to assess the health and density of vegetation on the Earth's surface, calculated using satellite imagery. Terra/MODIS 16 days 250m/500m Vegetation Index (Didan, 2021) was used for the nowcasting exercise.

parameters σ_u and σ_v based on both the standard deviation of the response variable and user input. The size and range of the grid can be expanded or narrowed as needed during the process based on model performance.

For the purpose of nowcasting, the predictive performance of one-period-ahead forecasts is used for the assessment of each model specifications. The performance is evaluated using the out-of-sample RMSE, calculated with one-period-ahead prediction errors for an expanding window of *N* quarters. The length of the expanding window, *N,* can be determined based on the out-of-sample and in-sample performance of the model specifications. An excessively long testing window may overly weigh periods with less relevance to the period of interest, while a window that is too short may result in misinterpreting spurious relationships as indications of strong predictive power.

It is also important to note that while the model selection is done based on predictive performance over the test period, the model estimation is done using the data throughout the sample period. Given a dataset with length *T*, the model generates one-period-ahead predictions at period $t = T - n + 1$ for each $n = N, \dots, 1$, using actual data available up to period t. Since the model produces a distribution of predictions rather than a single point of prediction for each period, the mean prediction error of each prediction is used to calculate the out-of-sample RMSE.

IV. Application to nowcasting Tanzania's growth

This section applies the nowcasting methodology to the case of Tanzania's economic growth.

Data Description

For the application to nowcasting of Tanzania's growth, monthly and quarterly economic indicators related to growth are added as high-frequency regressors.

The National Bureau of Statistics (NBS) of Tanzania releases quarterly GDP with a lag. Figure 2 summarizes year-over-year (yoy) quarterly growth of the economy over the past decade[7], highlighting the trends and fluctuations in the country's economic growth. Tanzania is one of Africa's fastest growing economies in the 2020s (Regional Economic Outlook: Sub-Saharan Africa), and a simple time series decomposition assuming additive structure shows fluctuations in real growth trend until recent years and a 'random' growth factor that remains unexplained by the assumed trend and seasonality (Figure 3). A better understanding of these dynamics through various economic indicators would be beneficial for policy decisions, as it can provide insights into the underlying drivers of economic performance.

Before the official quarterly GDP is published, a set of high-frequency indicators relevant to the country's economic activities becomes available. Agriculture, construction, and manufacturing sectors have been significant contributors to Tanzania's growth over the past decade. Agriculture, for instance, accounted for about 27 percent of GDP in 2023 and generated about 65 percent of total employment in 2021. Several

[7] Log level of real GDP is used as the variable of interest. For better comparability and interpretability, nowcast estimates reported were converted to year over year growth in the output tables. See Kohns and Bhattacharjee (2023) for another approach to address the issue of boundless drift with using growth variables.

available high-frequency indicators, including monthly rainfall data, the agricultural stress index, and the amount of cash crop exports, are closely linked to the performance of agricultural activities.

Industry and construction contributed approximately 30 percent of GDP in 2023. Notably, construction led the pre-pandemic growth with robust double-digit growth between 2016 and 2019. Services accounted for 39 percent of GDP in 2022, with the country's tourism and financial sector exhibiting robust growth post-pandemic. Available high-frequency indicators that are relevant to each major economic activity, such as industrial raw materials imports, service receipts, and credit to private sector, are included in the set of regressors. Additional monetary, fiscal, exports, and imports-related indicators are also included to capture macroeconomic movements in the economy. Table A 1 lists the available high-frequency predictors from various sources.

In operational real-time settings, nowcasting requires dealing with publication lags and data revisions. Because complete real-time vintages for Tanzania are not available, the model evaluation is based on pseudo-real-time data. High-frequency indicators are aligned to the quarterly reference period using simple aggregation or lags, and missing data at the edge of the sample are handled through standard transformations prior to model estimation (Table A 1). In an operational setting, users can interpret nowcasts with attention to data release calendars, monitor widening uncertainty bands as releases are delayed or revised, and treat shifts in inclusion probabilities or fit quality as early signals that indicator relationships may be evolving.

Figure 2. Tanzania: Real GDP growth (year-over-year, percentage)[8]

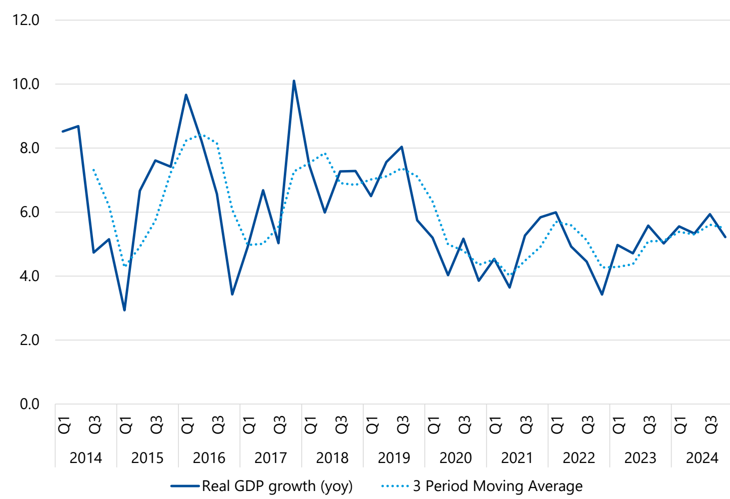

Source: Tanzania National Bureau of Statistics and staff calculations.

[8] The GDP growth series for Tanzania exhibits a decline in variance, particularly after the Covid-19 pandemic. BSTS can capture periods of lower volatility through its state-space structure and variance priors but does not model variance as a dynamic process in the way GARCH-type models do. While this is generally adequate for short-term forecasts, GARCH-type models can serve as a complementary benchmark for analyzing volatility patterns.

Figure 3. Tanzania: Real GDP growth Decomposition

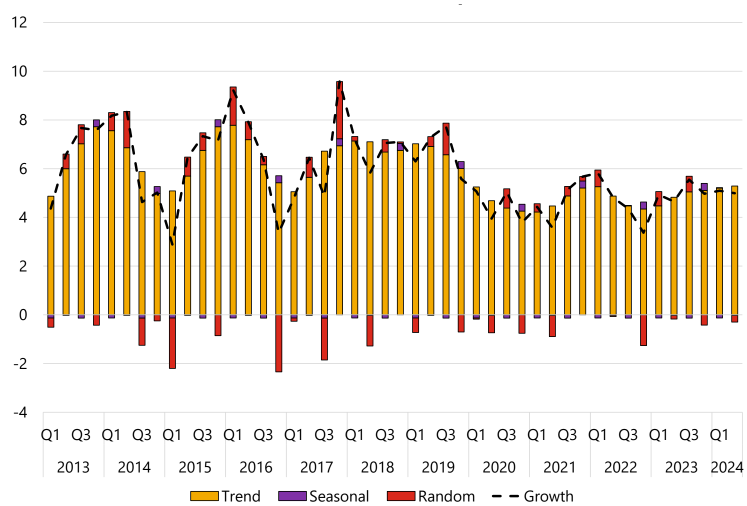

Source: Tanzania National Bureau of Statistics and staff calculations.

Note: GDP growth decomposition shows results from decomposing the growth series into seasonal, trend, and random components using moving averages, assuming additive seasonal component.

Application and Model Results

The set of economic indicators and model inputs described in Section 3 constitutes the model assumptions. Figure 4 illustrates the average errors and one-period-ahead predictions from the model specifications with the lowest out-of-sample RMSE. The estimates were calculated over an expanding test window of 16 quarters, spanning from 2021Q1 to 2024Q4, with the objective of nowcasting growth in 2025Q1. The test period represents roughly 25 percent of the full sample from 2008Q1 onward. This window length was chosen to ensure that the rankings of competing specifications are not driven by short-lived or spurious relationships, yet to focus on the period most relevant for operational nowcasting. For each quarter in the test window, the model is re-estimated using all information available up to that point, and a one-step ahead forecast is generated. To assess robustness, shorter and longer evaluation windows were also examined, with results presented in Figure 5.

Table 1 summarizes the performance and posterior parameter distributions of the model specifications with the lowest out-of-sample RMSE. Both naïve and AR(1) models, as well as 5-year quarterly averages, were used as benchmarks to assess relative performance.[9] The RMSE of averaged estimates over the 16-quarter test window was 0.46, which compares to the RMSE of 0.8 of a benchmark AR(1) model. To formally compare forecast accuracy, Diebold-Mariano (DM) (Diebold and Mariano (1995)) tests were conducted; the DM p-value for the BSTS model was 0.009 versus the AR(1) benchmark and 0.002 versus the naïve model, indicating statistical significance in predictive performance. The estimation precision improves toward the end of the sample period, largely explained by the expanding window approach allowing the model to incorporate a longer history of data as well as potential improvement of the quality of data over time.

[9] Simple benchmarks such as the AR(1) or naïve model approaches are used in nowcasting applications for their ease of implementation and long-standing role in macroeconomic analysis. For instance, Giannone et al. (2013) employed AR(1) as a benchmark for nowcasting China's real GDP, while Bok et al. (2017) adopted a naïve AR model in their study on nowcasting with big data for the United States. Qiu et al. (2018) and Kohns and Bhattacharjee (2023) present performance of BSTS model compared to other alternative models.

It is noteworthy that certain periods exhibit relatively abrupt changes in growth, which may have affected the relative performance of some model specifications. Real GDP growth in 2023 Q1, for example, showed a jump from 3.4 percent growth in the previous quarter to 5 percent, partly reflecting the pronounced contraction of crop production during 2022 Q4. While the model specifications that allow larger variations in trend components perform better at capturing the volatility, the models do not necessarily perform better throughout the test period. To assess the extent to which outliers are influencing the results in the nowcast, a robustness check was performed by excluding the data inputs from the volatile period. The top panel of Figure 5 summarizes the results from an alternative setup where the two periods are removed from the model, indicating that the outliers have not affected the nowcast results significantly.

To further assess the robustness of the BSTS model in GDP nowcasting, a sensitivity analysis was performed by varying key prior parameters:

- Trend variance (level, slope): Controls the flexibility of the trend component.
- Seasonal variance: Determines the level of fluctuation in seasonal effect.
- Expected model size: Reflects prior belief about the number of relevant regressors.

In economic terms, the priors that matter most are those governing the level and slope variances of the trend, as these control how flexibly the underlying growth path can adjust. Very tight priors make the trend overly rigid and suppress genuine turning points, while overly loose priors introduce excess noise. The sensitivity analysis suggests that calibrating the standard deviation parameters to roughly 2 to 5 percent of the standard deviation of the observed series yields stable results for Tanzania. However, the calibration should be data-dependent and should be adjusted in settings where the underlying trend exhibits greater intrinsic volatility. In particular, higher trend volatility or more frequent structural changes would warrant looser priors to ensure that the model remains responsive to the shifts. A practical approach would be to begin with a broad grid of potential value of parameters and refine it based on model performance and sensitivity results for the specific series of interest.

The seasonal component plays a role in capturing recurring intra-year patterns in economic activity, such as agricultural cycles, fiscal execution, and tourism-related fluctuations, while the expected model size prior mainly regulates the inclusion of noisy or weakly informative indicators. The model size prior encodes a prior belief about how many indicators are likely to be informative in the context of nowcasting. A smaller expected model size assumes that only a subset of indicators have significant predictive content. In this application, an expected model size of roughly 50-60 percent of the indicator set worked well for near-term forecasting, with scope to adjust the assumption as the sample expands and data quality improves.

The bottom panel of Figure 5 summarizes further detail how changes in these parameters affect forecast accuracy and the distribution of nowcasted estimates, using forecast distributions derived from the parameter combinations that achieved the lowest RMSE among all tested settings. Across the four sensitivity charts, forecast performance is most responsive to variance priors in the trend slope, with moderate values delivering the lowest out-of-sample RMSE; both overly restrictive and diffuse settings reduce accuracy. Increasing the prior model size yields a modest RMSE improvement but the result also highlights diminishing returns from adding predictors beyond a certain threshold with added noise. Overall, the results indicate relatively stable forecasts once priors fall within the moderate range and offer a useful benchmark for setting prior assumptions.

Figure 4. One-period-ahead Mean Predictions, Real GDP Growth of Tanzania (year-over-year, percentage)

One-period-ahead Prediction Errors, Averaged and Benchmark

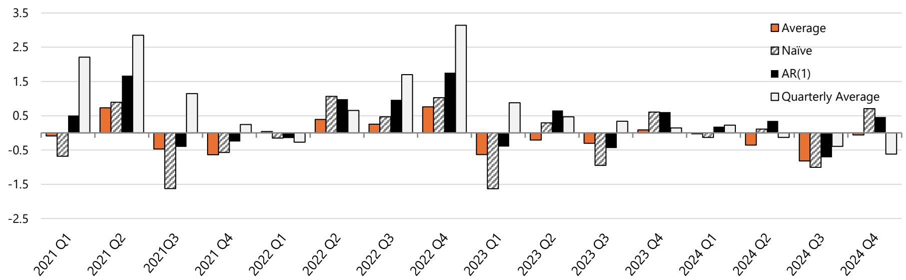

Distribution of One-period-ahead Mean Predictions, Model Specifications with Lowest Out-of-sample RMSE

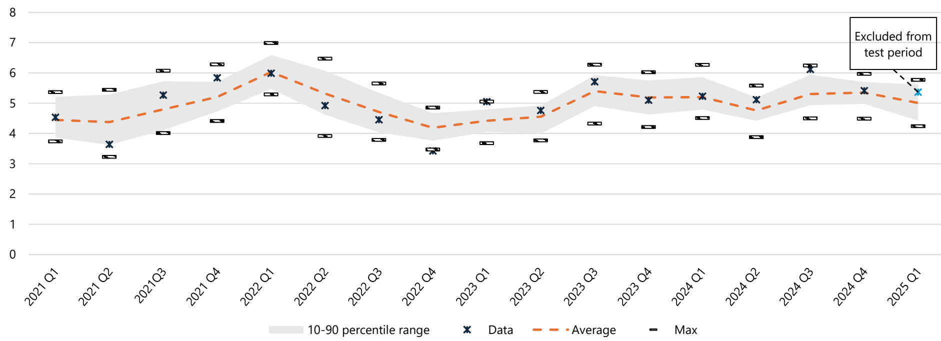

Note: The charts present results from twenty model specifications with lowest out-of-sample RMSE.

Table 1. Prior and Posterior Distribution of Model Parameters and In-sample Prediction Errors

		PRIOR GUESS				POSTERIOR DISTRIBUTION				IN-SAMPLE RESULTS	
RMSE	RMSE	Model Size	σ_μ (level)	σ_δ (slope)	σ_ω (season)	Model Size	σ_μ	σ_δ	σ_ω	Prediction Errors	
	Norm.						Mean			Mean	SD
0.42	0.53	14	0.021	0.003	0.03	3.1	0.015	0.003	0.018	0.015	0.020
0.49	0.61	14	0.027	0.003	0.02	3.5	0.018	0.003	0.013	0.017	0.026
0.56	0.70	18	0.006	0.009	0.06	5.1	0.006	0.007	0.033	0.015	0.019
0.57	0.71	14	0.009	0.006	0.06	3.1	0.008	0.005	0.033	0.017	0.022
0.61	0.76	18	0.021	0.006	0.05	5.2	0.017	0.005	0.028	0.014	0.018
0.63	0.79	20	0.006	0.006	0.03	6.0	0.006	0.005	0.018	0.014	0.019
0.63	0.79	16	0.003	0.012	0.05	4.2	0.003	0.009	0.028	0.016	0.022
0.64	0.79	16	0.012	0.012	0.04	4.2	0.011	0.009	0.023	0.015	0.021
0.65	0.81	10	0.006	0.006	0.02	1.7	0.006	0.005	0.013	0.014	0.019
0.65	0.81	10	0.012	0.012	0.05	2.1	0.011	0.009	0.028	0.015	0.020

Note: This table reports the parameter settings and posterior distributions for the model specifications generating the lowest out of sample RMSE for the forecasts, highlighting how forecast performance varies across alternative trend and variance priors. The second column presents normalized RMSE values using the benchmark AR(1) model RMSE.

Figure 5. Robustness checks

One-period-ahead Mean Predictions, Omitting Outliers during the Test Period

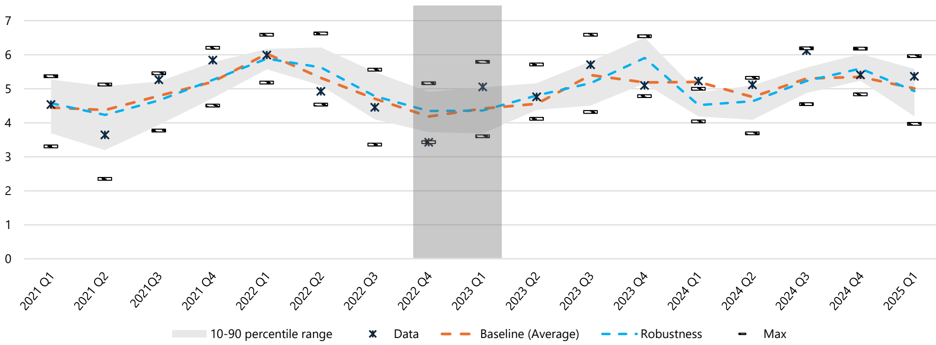

Sensitivity Analysis: Average Nowcasted 2025 Q1 Growth by Prior Values and Test Window Length

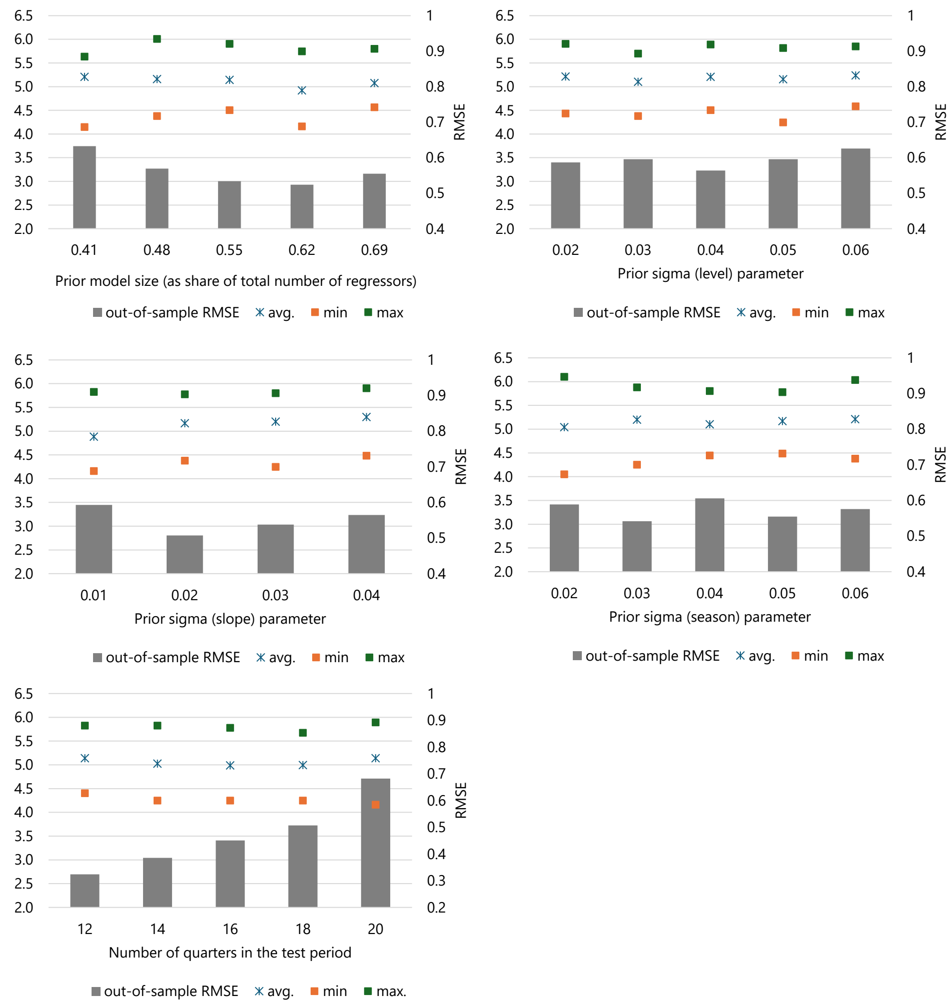

Note: The top panel presents the results from twenty model specifications with lowest out-of-sample RMSE. The bottom panel summarizes how different variance and model-size priors affect nowcast distributions and forecast errors. For easier applicability, the sigma parameter values are expressed as a share of standard deviation of the growth series (y). For operational use, a similar exercise would help identify “safe zones” where the output is stable and informs users on parameter settings where forecasts become sensitive, which could be useful for identifying when recalibration may be required.

Additional outputs from the model are useful for further understanding the growth patterns of the economy. Figure 6 presents the average contribution of trend, seasonality, and regression components to the predictions during the estimation stage. The trend component is the major contributor to growth forecasts, reflecting the model specification in log-level GDP. The pace of growth has been on average slower following the Covid-19 pandemic, which is captured primarily through adjustments in the underlying trend. The size of contributions from the seasonality component has largely remained stable, although there has been a change in patterns since 2020, with more frequent peaks in growth observed in Q3 rather than in Q4.

Regressors improve forecast accuracy by refining the latent trend and explaining short-run deviation around the path. The contribution from regression factors has fluctuated over the period, with only marginal contribution following the recovery from Covid pandemic. The inclusion probability of variables provides useful insights into the relative contribution of each regressor to the forecast, after accounting for seasonality and trend factors. Figure 7 shows the average posterior inclusion probability of the predictors that were assigned highest probabilities in the model specifications with lowest RMSEs. External demand indicators dominate, with Emerging Markets PMI (composite and manufacturing) and Developing Markets PMI showing the highest inclusion probabilities with each exceeding 40 percent, highlighting the association of global conditions with Tanzania's growth. Monetary aggregates such as M3 and Monetary Base, along with oil prices and private credit, also feature prominently, reflecting domestic liquidity and financial conditions. Agricultural Stress Index and commodity prices appear with moderate probabilities, consistent with the episodic nature of sector-specific shocks, particularly in agriculture, which remains sensitive to weather fluctuations and international commodity markets.

Collectively, these results underscore the role of external demand, monetary conditions, and sector-specific factors in shaping near-term growth dynamics. While high inclusion probabilities do not establish causality between the indicators and observed growth outcomes, they are broadly consistent with developments observed in Tanzania's economy, including a strong contribution from the agriculture as well as a broadening of growth drivers. The change in inclusion probabilities of regressors over the test period (Figure 8) indicates a general decline in the importance of most regressors, reflected in a smaller regression contribution, with the exception of external indicators, which remain robust. The temporary decline in the inclusion probability of external demand indicators, including the emerging markets PMI, during the Covid-pandemic suggests a temporary breakdown in the predictive relationship between global activity indicators and domestic growth during the pandemic, followed by a re-establishment of external demand as a key contributor as global conditions normalized. While the persistent selection of domestic fiscal indicator highlights the importance of fiscal conditions for near-term activity, the gradual decline over time indicates potential weakening of contemporaneous fiscal-growth links with growth drivers gradually broadening beyond the private sector toward construction, manufacturing, as well as services and tourism.

Figure 6. Contribution by Components

Mean Contribution to Predictions by Model Components, Average

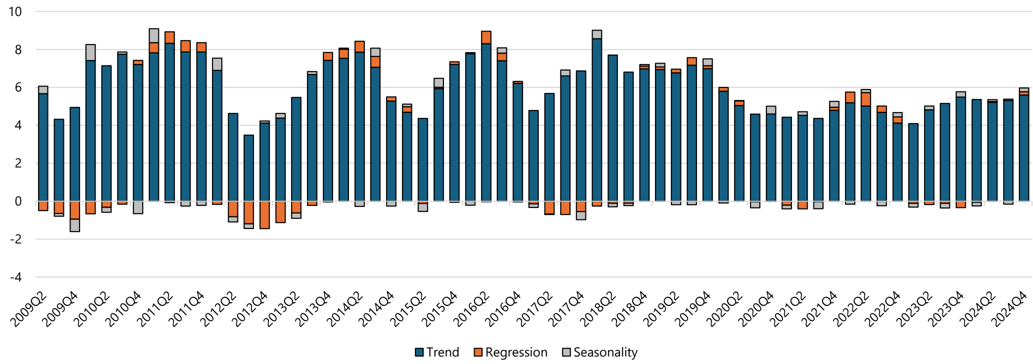

Seasonality, by Quarter

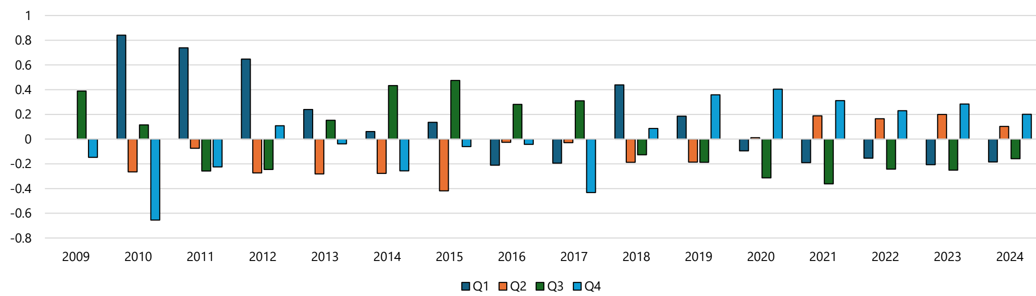

Note: The dominance of the trend component reflects the fact that the model is estimated in log GDP levels, where a stochastic trend captures most of the persistent variation in the series and regressors explain short-run deviations around this path, and lower volatility in the growth patterns post-2021. The resulting contribution was then converted to reflect contributions to resulting growth rate presented in the upper panel. Contribution by components averaged using twenty five model specifications with lowest out-of-sample RMSE for each case.

Figure 7. Inclusion Probability of Regressors

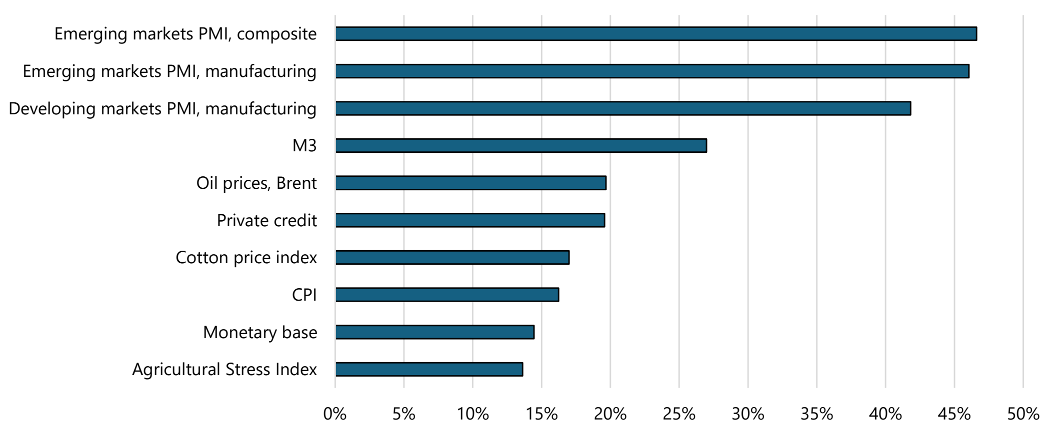

Note: The figure presents the inclusion probability of ten regressors with highest probabilities, averaged using twenty five model specifications with lowest out-of-sample RMSE.

Figure 8. Inclusion Probability over the Sample Period, Selected Indicators

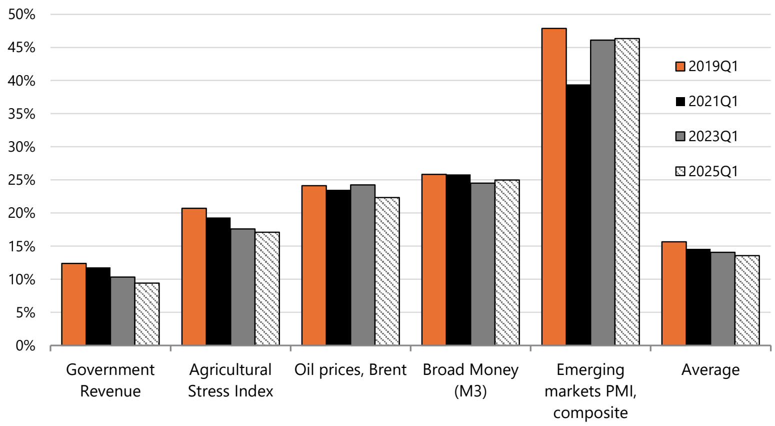

Note: Inclusion probabilities of the regressors were averaged using twenty five model specifications with lowest out-of-sample RMSE for each case.

V. Extensions

Extending the number of forecast horizons

The BSTS model can be extended to forecasting over multiple horizons. Tanzania's growth forecasting exercise illustrated in this section focuses primarily on near-term forecasts to ensure that the economic indicators included in the model remain relevant to growth during the forecasted period. However, the model can also be readily extended to produce long-term forecasts, primarily relying on trend and seasonality assumptions.

A similar model selection exercise is conducted as in the case of one-period-ahead nowcasting. A set of model assumptions is tested using an expanding window of *N* quarters, and the model generates *p*-period-ahead forecasts at the end of each training period. The out-of-sample RMSE is calculated using prediction errors where actual data is available using an expanding window approach similar to the one-period-ahead case. Specifically, the model generates forecasts for each quarter in the test window by using all available data up to that point, and then compares the predicted values to the actual observed GDP figures *p*-quarters ahead. The prediction errors for each forecast horizon are calculated as the difference between the forecasted and actual values to obtain the out-of-sample RMSE. This process is repeated for each model specification, and the nowcasts from specifications with the lowest RMSE are averaged.

Figure 9 shows the results for three-period-ahead forecasts with the model specifications with lowest out-of-sample RMSE. The test window of 16 quarters remained the same as in the baseline one-period-ahead case. Table 2 summarizes the prior- and post- model specifications and their performance. At longer forecast horizons, performance becomes increasingly sensitive to trend specification. While tighter trend variance priors can help filter short-term noise, they can also induce persistent bias when extrapolated over multiple period. Accordingly, the preferred multi-period specifications are characterized, on average, by higher trend variance priors compared to the baseline case. Consistent with this, the preferred multi-period specifications tend to feature smaller effective model sizes, indicating that the marginal contribution of additional regressors diminishes as the forecast horizon lengthens and the model relies more heavily on structural components such as the trend and seasonality.

Figure 9. Distribution of Mean Predictions, Three-period-ahead Forecasts

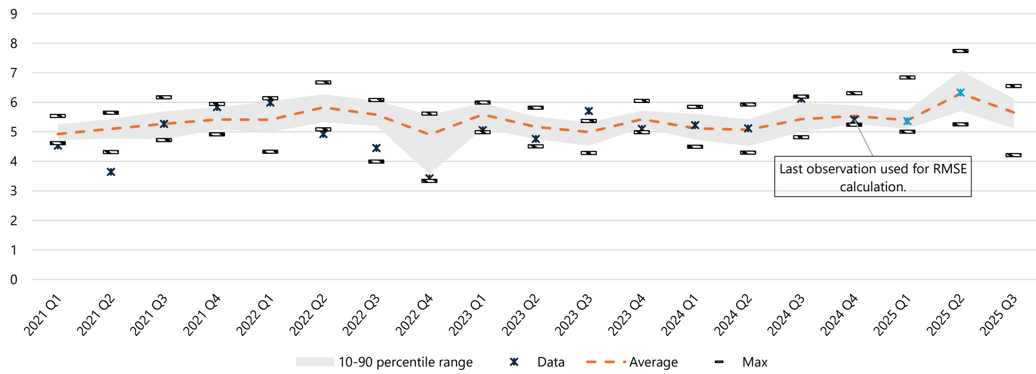

Note: The charts present results from twenty five model specifications with lowest out-of-sample RMSE.

Table 2. Prior and Posterior Distribution of Model Parameters and In-sample Prediction Errors, Three-period-ahead Forecasts

	PRIOR GUESS				POSTERIOR DISTRIBUTION				IN-SAMPLE RESULTS	
RMSE	Model size	σ_μ	σ_δ (slope)	σ_ω (season)	Model size	σ_μ	σ_δ	σ_ω	Prediction Errors	
						Mean			Mean	SD
0.77	4	0.075	0.014	0.04	1.4	0.048	0.003	0.027	0.017	0.023
0.79	4	0.067	0.011	0.03	0.5	0.042	0.003	0.021	0.015	0.020
0.79	6	0.006	0.003	0.03	1.1	0.006	0.007	0.018	0.016	0.021
0.80	2	0.019	0.003	0.03	0.2	0.015	0.005	0.019	0.015	0.020
0.80	2	0.075	0.008	0.05	0.6	0.048	0.005	0.032	0.015	0.019
0.80	4	0.083	0.011	0.05	0.6	0.053	0.005	0.033	0.015	0.020
0.80	2	0.028	0.006	0.05	0.3	0.021	0.009	0.030	0.015	0.020
0.81	4	0.050	0.011	0.05	0.4	0.035	0.009	0.031	0.015	0.020
0.81	2	0.039	0.008	0.02	0.2	0.026	0.005	0.014	0.015	0.020
0.81	4	0.028	0.006	0.05	0.4	0.021	0.009	0.030	0.015	0.020

Note: This table reports the parameter settings and posterior distributions for the model specifications generating the lowest out-of-sample RMSE for three-period-ahead forecasts.

Dynamic regressors and climate change

The BSTS model also allows for the inclusion of dynamic regressors. Correlation between certain economic indicators and growth may undergo change over the sample period, which may not be adequately captured by fixed regression coefficients. Sectoral transformations or external shocks may lead to such dynamic relationships. Incorporating dynamic regressors into the model can help address such potential changes in the relationship between an indicator and growth. Dynamic regression allows the coefficients of these dynamic regressors, $\beta_{i,t}$, to fluctuate over the sample period, as in the equations below. Unlike regular regressors, which are included based on estimated inclusion probabilities, dynamic regressors are always incorporated into the model regression.

$$\beta_{i,t+1} = \beta_{i,t} + \epsilon_t$$
$$\epsilon_t \sim N\left(0, \frac{\sigma_i^2}{variance_{x_i}}\right) \tag{4}$$

One such example is the correlation between rainfall and growth. The amount of rainfall can significantly impact growth through its effects on the agricultural sector and infrastructure, particularly in cases of floods and droughts. As a country undergoes sectoral transformation or experiences shifts in rainfall patterns – resulting in more frequent floods and droughts – the cumulative rainfall may exhibit a different association with growth. Figure 10 displays the changes in rainfall patterns in Tanzania over recent years, suggesting potential advantages of allowing the coefficient of the rainfall indicator to fluctuate over the sample period.

To examine the potential dynamic relationship, the model was re-estimated with rainfall included as a dynamic regressor, while other indicators remained static as in the *baseline* model. The test window between 2021 Q1 and 2024 Q4 was used to capture the periods with a large shift in the rainfall pattern. Figure 12 compares the in-sample forecast performance of the dynamic model with that of the baseline model presented in Section 4. The result suggests that the in-sample performance gain is mostly concentrated during the early sample period. The result is in line with the change in average coefficient value of the dynamic regressor (Figure 13), which shows that the estimated rainfall coefficient varies substantially over time. The negative coefficients during the earlier sample period coincide with documented episodes flooding, when rainfall acted primarily as a disruptive shock rather than a productivity input. The coefficient turns positive beginning 2016, consistent with rainfall supporting agricultural output under more normal conditions. From 2017 onward, the relationship weakens and becomes intermittently negative, reflecting structural transformation and the increasing volatility of rainfall patterns.

Figure 10. Tanzania: Rainfall Patterns
Trend in Annual Mean Precipitation (mm)

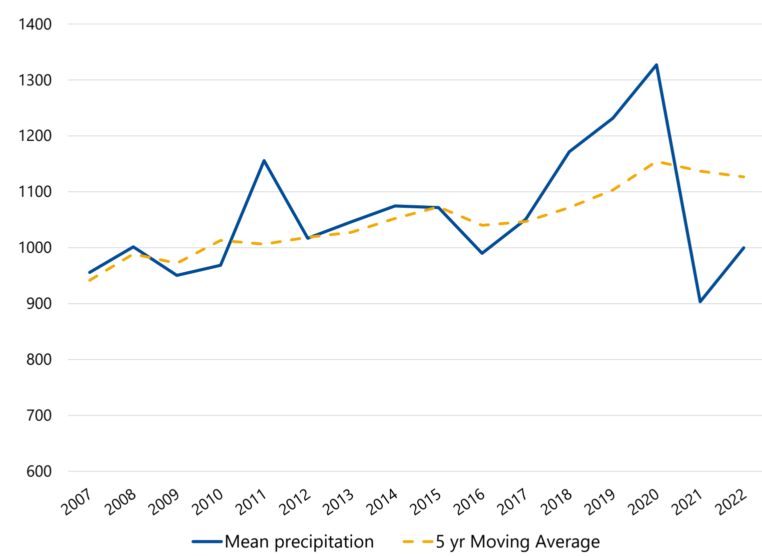

Source: World Bank Climate Change Knowledge Portal.

Trend in Number of Floods

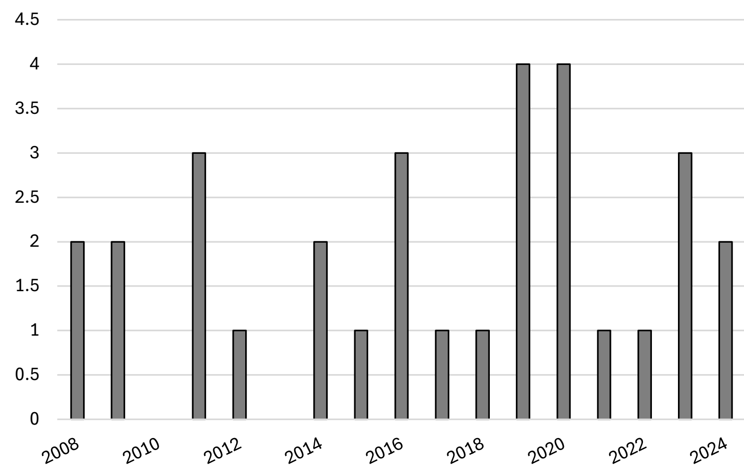

Source: World Bank Climate Change Knowledge Portal.

Figure 11. One-period Ahead Predictions, Dynamic Regressors

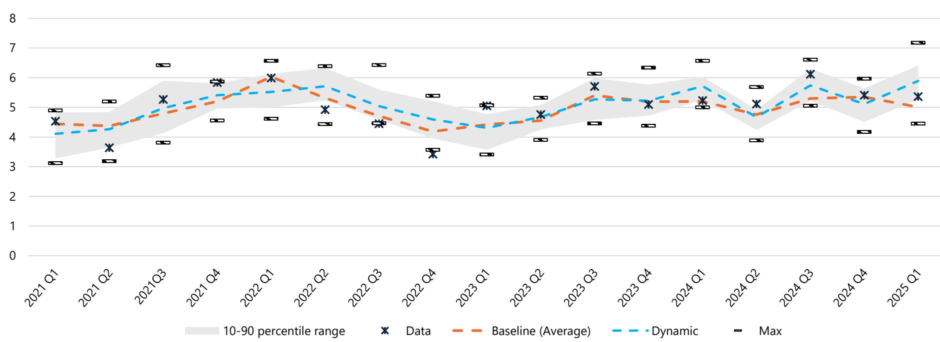

Note: The results shown in the charts are from the twenty five dynamic model specifications with the lowest RMSE, using Rainfall indicator as the dynamic regressor.

Figure 12. Cumulative Absolute One-period Ahead Forecast Errors, Percentage Points

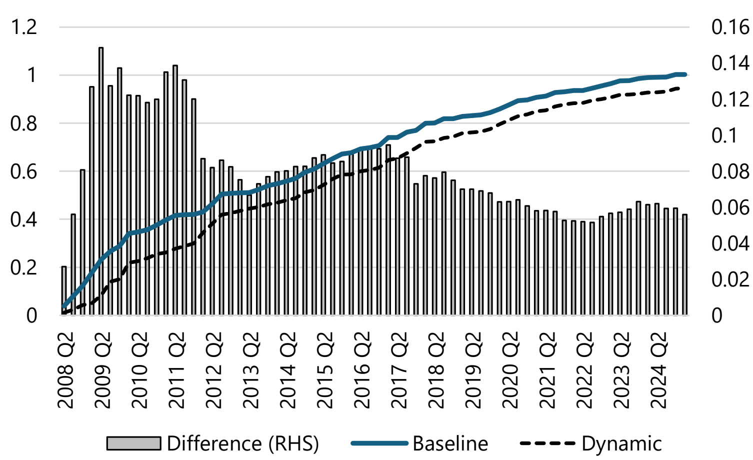

Note: The results shown in Figure 11 uses the results from the dynamic model specification with the lowest RMSE, using Rainfall indicator as the dynamic regressor.

Figure 13. Dynamic Coefficient over the Sample Period

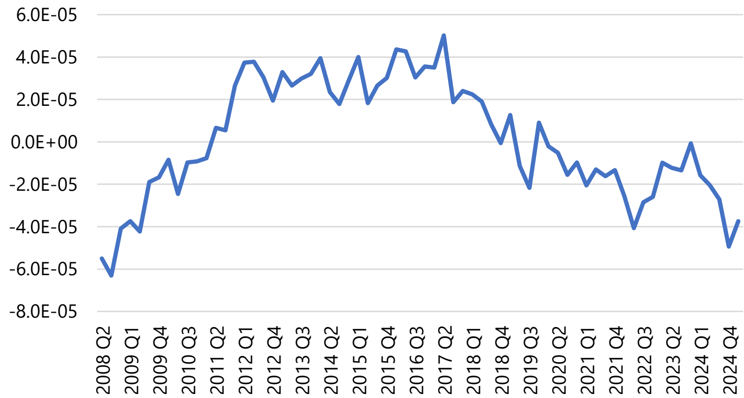

Note: Figure 12 shows the average dynamic coefficient values over the sample period, using the ten model specifications with the lowest RMSE. They illustrate how the dynamic indicator may become more (or less) informative as changes occur to its structural relationship with the response variable.

VI. Conclusion

This paper outlines a framework to apply the BSTS model to produce nowcast outputs for countries without well-established growth patterns, using high-frequency indicators. The BSTS model is a useful tool for nowcasting economic growth, particularly in the context of data gaps and volatility. The application to Tanzania's growth nowcasting shows that the approach produces reliable forecasts, outperforming the benchmark model. By allowing for flexible modeling of trends and seasonality, as well as the inclusion of dynamic regressors, the model can enhance the understanding of growth patterns and support real-time economic assessments. Prior beliefs and expert knowledge about the economy and growth can also be easily incorporated into the model as inputs for better performance and more reliable outputs. Furthermore, additional outputs from the model contribute to deeper insights into these growth dynamics.

Beyond its empirical performance, the BSTS framework offers several practical advantages for real-time monitoring in data-constrained environments. In operational settings, the model can be re-estimated on a regular basis once calibrated, incorporating newly released high-frequency indicators with a limited set of model specifications to be tested. The resulting predictive distribution provides policymakers with more than a point estimate. The posterior mean can serve as a central nowcast, while widening forecast bands or notable tail risks can signal rising uncertainty. Similarly, declining stability in posterior inclusion probabilities, or abrupt shifts in the relative importance of indicators, can indicate evolving macroeconomic relationships that warrant closer scrutiny. These diagnostic features, together with transparent decomposition of trend and indicator-level contributions, make the BSTS approach a practical, policy-relevant modeling tool.

At the same time, the results underscore the importance of careful model calibration. Forecast performance is sensitive to assumptions governing trend flexibility and variable selection. While the spike-and-slab prior provides a way to filter noisy indicators, it may under-select informative variables in short samples. These features highlight the need for robustness checks, sensitivity analysis, and informed prior calibration, all of which would be integral to the framework's practical application.

Despite these limitations, the BSTS setup remains broadly applicable to other low-income countries with similar data constraints. Its ability to accommodate growth data of varying quality, incorporate a wide range of high frequency indicators with uncertain predictive value, and provide transparent, component-level interpretations makes it a practical tool for real-time economic monitoring. Future research could explore these extensions, such as time-varying volatility or alternative shrinkage priors, while preserving the operational simplicity that makes the framework appealing for surveillance and policymaking settings.

Annex

Table A 1. High-frequency Economic Indicators

Variable (year-over-year growth, unless otherwise noted)	Source	Freq.	Description	Transformation
M3	BoT	M	M3	Moving Average
Monetary Base	BoT	M	Monetary Base, End Period	Moving Average
Reserve Money	BoT	M	Average Reserve Money	One-period lag
Private Credit	BoT	M	Credit to Non-government Sector	Moving Average
Exports of Goods	BoT	M	Exports of Goods	
Traditional Exports	BoT	M	Traditional Exports	Moving Average
Other Exports	BoT	M	Other Exports	Moving Average
Service Receipts	BoT	M	Service Receipts	Moving Average
Imports of Goods	BoT	M	Imports of Goods (f.o.b)	Moving Average
Capital goods imports	BoT	M	Capital goods imports	Moving Average
White products imports	BoT	M	White products imports	
Industrial imports	BoT	M	Fertilizers and industrial raw materials imports	
Food imports	BoT	M	Food and food stuffs imports	
Other imports	BoT	M	All other consumer goods imports	Moving Average
Service imports	BoT	M	Imports of Services	One-period lag
CPI	BoT	M	CPI	One-period lag
Rainfall	BoT	M	Rainfall (mm)	
ASI[1/]	FAO	M	Agricultural Stress Index	Moving Average
Oil prices, Brent	IMF	M	Crude Oil (petroleum), Dated Brent, light blend 38 API, fob U.K., US$ per barrel	One-period lag
Cotton price index	IMF	M	Cotton, Cotton Outlook 'A Index', Middling 1-3/32 inch staple, CIF Liverpool, US cents per pound	
Revenue	MoF	M	Gov. Total Revenue	Moving Average
VAT revenue	MoF	M	VAT on non-petroleum imports	One-period lag
Expenditure	MoF	M	Total Expenditure	Moving Average
Development Expenditure	MoF	M	Development Expenditure	Moving Average
Wage expenditure	MoF	M	Wages and Salaries	
Tourist arrivals	NBS	M	Tourist Arrivals	
Electricity	NBS	M	Electricity Consumption	
Cement	NBS	Q	Cement, manufactured	Moving Average
Emerging markets PMI, manufacturing	S&P Global	M	Emerging Markets PMI: Manufacturing (SA, 50+=Expansion)	Moving Average
Developed markets PMI, manufacturing	S&P Global	M	Developed Markets PMI: Manufacturing (SA, 50+=Expansion)	Moving Average
Emerging markets PMI, composite	S&P Global	M	Emerging Markets PMI: Composite (SA, 50+=Expansion)	Moving Average
Covid	IMF	M	Dummy variable (=1 for periods between 2020M1 and 2022M12)	

Note: BoT refers to Bank of Tanzania; MoF refers to Tanzania Ministry of Finance; NBS refers to the National Bureau of Statistics of Tanzania; FAO refers to Food and Agriculture Organization of the United Nations.
[1/] The level of ASI was used for the nowcasting exercise.

References

Akbal, O., Choi, S., Narita, F., and Yao, J. 2023. “Panel nowcasting for countries whose quarterly GDPs are unavailable.” IMF Working Paper, 23/158.

Armesto, M. T., Engemann, K., and Owyang, M., 2010. “Forecasting with mixed frequencies.” Review, 92, issue Nov, p. 521-536.

Bok, B., Caratelli, D., Giannone, D., Sbordone, A. M., Tambalotti, A. 2018. “Macroeconomic Nowcasting and Forecasting with Big Data.” *Annual Review Economics*. 10:615-643.

Banbura, M., Giannone, D., and Reichlin, L., 2011. Nowcasting. Oxford Handbook on Economic Forecasting.

Barhoumi, K., Choi, S., 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.” IMF Working Paper, 22/88.

Bear, A., Guerreiro, V., and Silungwe, A., 2022. “2020 Global Stocktaking of National Accounts Statistics: Availability for Policy and Surveillance” IMF Working Paper No. 22/29.

Clements, M. P., and Galvão, A. B., 2008. “Macroeconomic Forecasting with Mixed-Frequency Data: Forecasting Output Growth in the United States.” *Journal of Business & Economic Statistics*, *26*(4), 546–554. http://www.jstor.org/stable/27639011

D'Agostino, A. and Gambetti, L. and Giannone, D., 2010. “Macroeconomic Forecasting and Structural Change.” ECB Working Paper No. 1167.

Diebold, F. and Mariano, R., 1995. “Comparing Predictive Accuracy” Journal of Business & Economic Statistics, July 1995, Vol. 13, No. 3.

Didan, K., 2021. *MODIS/Terra Vegetation Indices 16-Day L3 Global 500m SIN Grid V061.* distributed by NASA EOSDIS Land Processes Distributed Active Archive Center, https://doi.org/10.5067/MODIS/MOD13A1.061. Date Accessed 2024-12-01.

Giannone, D., Reichlin, L., and Small, D., 2008. "Nowcasting: The real-time informational content of macroeconomic data," Journal of Monetary Economics, Elsevier, vol. 55(4), pages 665-676, May.

Giannone, D., Agrippino, S. M., & Modugno, M., 2013. Nowcasting China Real GDP. CIRANO.

International Monetary Fund. 2024. *Regional Economic Outlook: Sub-Saharan Africa— A Tepid and Pricey Recovery.* Washington, DC, April.

Kohns, D., and Bhattacharjee, A., 2023. "Nowcasting growth using Google Trends data: A Bayesian Structural Time Series model," International Journal of Forecasting, Elsevier, vol. 39(3), pages 1384-1412.

Mariano, R. and Murasawa, Y., 2010. A Coincident Index, Common Factors, and Monthly Real GDP*, Oxford Bulletin of Economics and Statistics, 72, issue 1, p. 27-46.

Pérez, F., 2018. "Nowcasting Peruvian GDP using Leading Indicators and Bayesian Variable Selection," Working Papers 2018-010, Banco Central de Reserva del Perú.

Scott, S. L., 2024. bsts: Bayesian Structural Time Series. R package version 0.9.10, <https://CRAN.R-project.org/package=bsts>.

Scott, S. L. and Varian, H. R. 2014. Predicting the present with Bayesian structural time series. International Journal of Mathematical Modelling and Numerical Optimisation 5, 4–23.

Scott, S. L. and Varian, H. R. 2015. Bayesian variable selection for nowcasting economic time series. In A. Goldfarb, S. Greenstein, and C. Tucker, eds., Economics of Digitization, 119 –136. NBER Press, London.

Qiu, J., Jammalamadaka, S. R., and Ning, N. 2018. Multivariate Bayesian Structural Time Series Model. Journal of Machine Learning Research 19 (2018) 1-33.

Woloszko, N., 2020. "Tracking activity in real time with Google Trends", OECD Economics Department Working Papers, No. 1634, OECD Publishing, Paris, https://doi.org/10.1787/6b9c7518-en.

World Bank, Climate Change Knowledge Portal, 2024. URL: https://climateknowledgeportal.worldbank.org/. Date Accessed: 2024-12-01.

PUBLICATIONS

Nowcasting Growth Using the Bayesian Structural Time Series Model: Application to Tanzania
Working Paper No. WP/2026/049