Improving Existing Methods of IFSA Supply Forecasting

by

Walter Ac-Pangan, Nathan P. Hendricks, Yacob Zereyesus, Jennifer Kee, Jeremy Jelliffe, Stephen Morgan, Lila Cardell, and Noé J. Nava

Improving existing methods of IFSA supply forecasting

Walter Ac-Pangan, Nathan P. Hendricks[1], Yacob Zereyesus, Jennifer Kee, Jeremy Jelliffe, Stephen Morgan, Lila Cardell, Noé J. Nava.[2]

Paper presented at the NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management
St. Louis, Missouri, April 24-25, 2023.

Disclaimer: The findings and conclusions in this presentation are those of the authors and should not be construed to represent any official USDA or U.S. Government determination or policy. This research was supported in part by the U.S. Department of Agriculture, ERS.

[1] Walter Ac Pangan is a Graduate Research Assistant and Nathan Hendricks is a Professor in the Department of Agricultural Economics at Kansas State University.
[2] Yacob Zereyesus, Jennifer Kee, Jeremy Jelliffe, Stephen Morgan, Lila Cardell, Noé J. Nava are Research Agricultural Economists at USDA Economics Research Services.

Improving existing methods of IFSA supply forecasting

The International Food Security Assessment (IFSA) report provides forecasts for grain demand, production, and the implied additional grain supply requirement for 77 low- and middle-income countries. In this study, we attempt to enhance the IFSA model by improving the model to forecast production 1 and 10 years into the future. Our results indicate that forecasting growing area and yield separately performs the best when forecasting production rather than directly forecasting production. For predicting and forecasting growing areas, the best model specification includes country-specific linear trends, annual precipitation, futures price, and country-specific fixed effects. Similarly, the best model specification for predicting and forecasting yield involves country-specific linear trends, pooled coefficients on temperature, and country-specific fixed effects. When we compare our best model specification with previous methods used in the IFSA model to predict yield, our revised model outperforms the previous method.

Key words: Global supply, forecast, food security

Introduction

The International Food Security Assessment (IFSA) is an annual report from the USDA-ERS that provides a 1 and 10-year projection of food security indicators for 77 low- and middle-income countries. The IFSA report forecasts demand for grains (food and non-food), grain production, and the implied additional grain supply requirement. The implied additional grain supply requirement quantifies the total grain demand in each region and sub-region that is not projected to be met through domestic production (Baquedano et al., 2021). Estimating the implied additional grain supply requirement is important for assessing international trade opportunities and potential areas of food aid needs.

Increasing population, per capita income, climate change and shortfalls in outputs are factors that increase the gap between demand and production. Growing demand for agricultural commodities from regions and countries with increasing populations and per capita income, such as Southern Africa and Asia, can open potential international trade opportunities and markets for U.S. agricultural products.

Our study attempts to improve the IFSA production forecasts in four ways. First, we allow models to select either pooled or country-specific coefficients. Pooled coefficients reduce the number of parameters that must be estimated, while country-specific coefficients allow the model to capture more heterogeneity. Second, we incorporate weather and price variables into the forecasting models to capture the influence of these variables on grain production. Third, we forecast growing area and not just yield. Finally, we select the best model specification using a time-series cross validation technique that accounts for time series nature of the data and the challenge of forecasting in the future. This study focuses on finding the model that provides the best out-of-sample forecast ten years into the future.

Forecasting production is a challenge and becomes especially complex when forecasting far into the future. A previous study evaluated the IFSA models using a cross-validation approach to select the best model specification for forecasting. They found that disaggregating yield estimates at the subregional level improved yield prediction (Zereyesus et al., 2022). Our study builds on this previous work. Another related study is Haile et al. (2016) who perform an out-of-sample forecast of crop acreage in 2014.

Our work also builds on several other strands of literature that model global grain production. First, there is large literature that estimates supply response using futures prices as explanatory variables (Haile et al., 2016; Hendricks et al., 2014; Roberts & Schlenker, 2013). A second strand of literature studies how climate change is likely to affect global production. For example, Lobell et al. (2011) use a panel of country-level data to estimate a model of yield that includes temperature, precipitation, and country-specific fixed effects. A third strand of literature estimates heterogeneous—and potentially nonlinear—yield trends across countries to forecast future global supply of food (Grassini et al., 2013; Ray et al., 2012).

Methodology

The objective of our modeling is to find the specification that is best at out-of-sample forecasting of grain production in the 77 IFSA countries. To accomplish this objective, we establish a general specification and then evaluate several different variations of the general specification. We estimate out-of-sample forecasting metrics for each specification and select the one that performs the best. Note that throughout this section, we consider aggregate grain production rather than production of individual crops.

Direct versus Indirect Production Forecasts

Grain production has two components: growing area and yield (i.e., production per unit of area). This is written as:

$$Prod_{it} = Area_{it} \times Yield_{it} \qquad (1)$$

where $Prod_{it}$ is the production in country i and year t, $Area_{it}$ is the growing area, and $Yield_{it}$ is the yield.

We consider a direct production forecast and an indirect production forecast. A direct production forecast estimates a model to directly forecast production ($Prod_{it}$) and does not create any forecasts of area or yield. An indirect production forecast estimates a model to forecast growing area ($Area_{it}$) and a separate model to forecast yield ($Yield_{it}$). Then it creates a forecast of production by multiplying the forecasted area times the forecasted yield.

Econometric Specifications

The general econometric specification to forecast production is written as:

$$Y_{it} = f(\theta_i, t) + \beta_i X_{it} + \alpha_i + \epsilon_{it}, \qquad (2)$$

where Y_{it} denotes the dependent variable: production (Millions of Metric Tons), area (Hectares), or yield (Millions of Metric Tons per Hectare); $f(\theta_i, t)$ denotes a trend that could either be linear or nonlinear; β_i are the coefficients for each independent variable; X_{it} are the independent variables (price, temperature and precipitation); α_i are country fixed effects; and ϵ_{it} are the residuals. In the most general specification, we allow country-specific trends and country-specific coefficients on predictors so that the coefficients have an *i* subscript (θ_i and β_i). It is important to note that our objective is to forecast production into the future and not to provide causal estimates of how different factors affect production. Therefore, we do not impose economic theory on the specification. Instead, we keep the specification as flexible as possible and allow out-of-sample prediction errors to guide the model selection. Next, we summarize each of the variations on the general model specification in equation (2) that we considered and assess the forecast accuracy of each specification.

Linear versus Poisson

We estimate equation (2) as a linear model using ordinary least squares (OLS) regression. However, we also estimate the model as a Poisson regression. Poisson models the function as $Y_{it} = \exp(\cdot)$, which is similar to an OLS model where the dependent variable is a logarithm (i.e., $\ln(Y_{it})$). The linear specification assumes that a change in the predictor affects the dependent variable by a certain number of units. The Poisson specification assumes that a change in the predictor affect the dependent variable by a certain percentage.

Linear versus Nonlinear Trends

The trend function is either assumed to be linear or a flexible nonlinear function. When a nonlinear relationship is allowed, we use a natural cubic spline with three knots. Natural cubic splines have several advantages over polynomials in that they can allow non-symmetric nonlinear trends with relatively few parameters estimated.

Pooled versus Country-Specific Trends

We consider the option where there is a single trend for all countries (i.e., pooled) or allow each country to have its own trend. The pooled model requires estimating fewer parameters, but country-specific trends are more flexible.

Various Controls

We consider specifications with different controls included in X_{it}. One option that we consider is to have no additional controls and model production simply according to trends. However, we also add different combinations of prices, precipitation, and temperature as predictors.

Pooled versus Country-Specific Coefficients on Predictors

Similar to the coefficients on trends, we estimate a single coefficient for each of the predictors (i.e., pooled) or we allow the coefficients to differ across countries.

Fixed Effects

The specifications also include some that have a common intercept across countries and others with country-specific intercepts (i.e., country fixed effects).

Time Series Cross-Validation

The accuracy of the model specifications is assessed using time series cross validation. Traditional cross validation trains the model (i.e., estimates the coefficients) with a portion of the data and then predicts the outcome on a different portion of the data that was not used for training. The data used to estimate the model is referred to as the training data and the data to assess the prediction accuracy is referred to as the test data. Cross-validation assesses the accuracy of the model at predicting out-of-sample. Time series cross validation is similar, except that it accounts for the time series nature of the forecast. For example, the model is trained with data up to a certain year and then the model is used to predict for a year in the future (Hyndman & Athanasopoulos, 2021).

Figure 1. *Out-sample forecasting models assessment.*

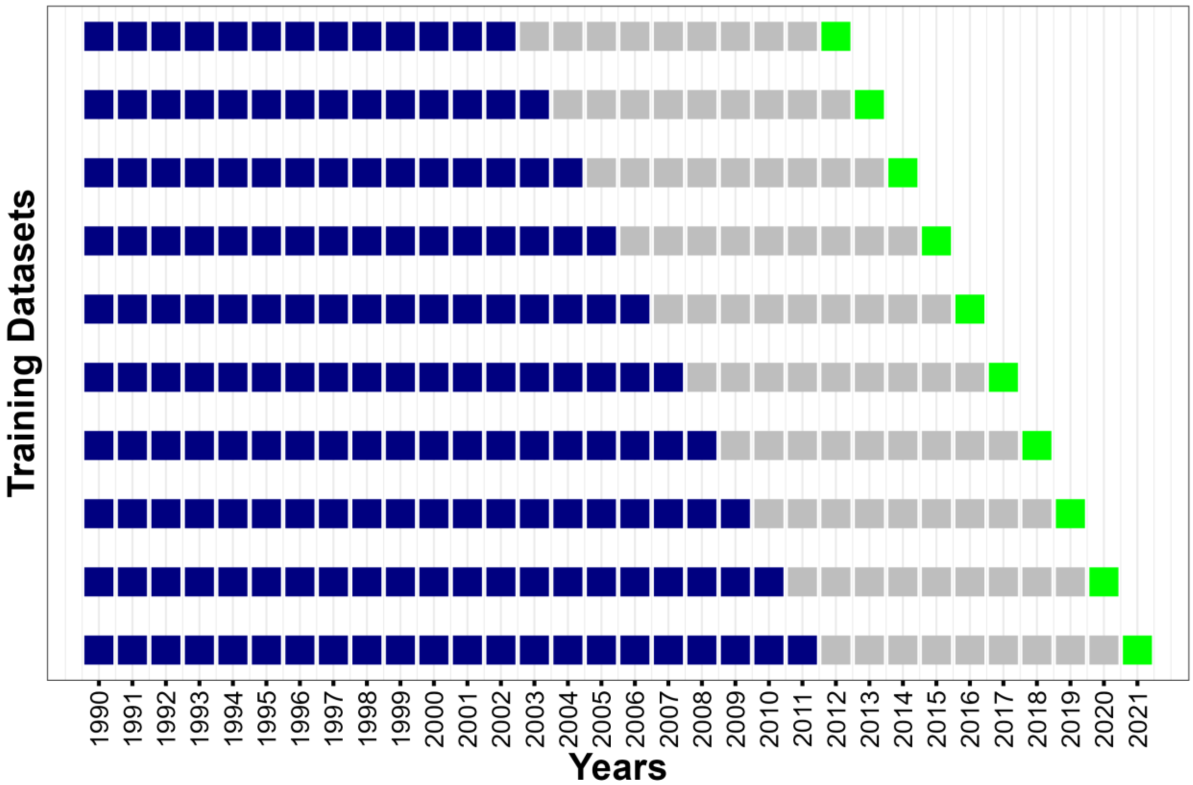

The time series cross-validation procedure that we used is illustrated in Figure 1. Our primary objective is to create an accurate forecast 10 years into the future. The data in our sample begins in 1990. First, we train the model with data between 1990 and 2002 and then predict the outcome in 2012. Next, we train the model with data between 1990 and 2003 and predict the outcome in 2013. We continue to repeat this process until we forecast the most recent year of observed data, 2021. This gives ten different estimates of the 10-year ahead prediction error for each specification that are summarized as described in the next section.

Accuracy Metrics

We estimate the difference between the out-sample predicted production values and the real production values present in the test dataset,

$$\widehat{\epsilon_{it}} = \widehat{Y_{it}} - Y_{it} \qquad (3)$$

where $\widehat{\epsilon_{it}}$ is the prediction error, $\widehat{Y_{it}}$ is the predicted value of the outcome, and Y_{it} is the observed value of the outcome in the test data. The prediction errors are then used to calculate the root mean squared error for each country,

$$RMSE_i = \sqrt{\frac{\sum_{t=\tau}^{T} \widehat{\epsilon_{it}}^2}{10}}, \qquad (4)$$

where τ is the earliest year in the test data and T is the final year in the test data. We make forecasts for 10 different years. We also calculate the mean absolute error for each country,

$$MAE_i = \sqrt{\frac{\sum_{t=2012}^{2021} |\widehat{\epsilon_{it}}|}{10}}. \qquad (5)$$

To create a single measure of the average predictive accuracy across all countries, we then calculate the average predictive accuracy as $\frac{1}{N}\sum_{i=1}^{N} RMSE_i$ and $\frac{1}{N}\sum_{i=1}^{N} MAE_i$, where N is the total number of countries. The preferred specification has the smallest RMSE or MAE. The RMSE and MAE both provide a valid metric of predictive accuracy and there is no clearly dominant metric. The RMSE places more weight on predictions that far away from the actual value by calculating the squared error. In this report, we will emphasize the RMSE results because we prefer to not be far off in a prediction for a country. However, we present results for both RMSE and MAE.

Forecasting

After we use the predictive metrics to select the preferred specification, then we use this specification to forecast production for 1 year and 10 years into the future (i.e., 2023 and 2033). To create a forecast of production in 2023 and 2033, we need to define values for any controls that are included in the preferred specification. We do not attempt to forecast prices or weather in the future, so we use the most recent 5-year average for each predictor as the assumed value for future years.

Data

Data were collected from 1960 until 2021 for as many of the variables as possible. However, some variables (World Price Index) were only available starting in 1990, so the training data begin in 1990. Grains included in our analysis include maize, rice, wheat, sorghum, barley, and millet.

Production, Area, and Yield

Production, growing area, and yield data were collected from the USDA Foreign Agricultural Service Production, Supply, and Distribution (PSD) database and from the UN Food and Agricultural Organization FAOSTAT database. We present results with the PSD data as our preferred specification. Production is measured in metric tons, area is measured in hectares, and yield as tons per hectare. Total grain production and area were aggregated by adding tons of production or hectares across crops. Yield was calculated as total grain production divided by area.

Prices

Three different measures of crop prices were considered in the specifications. One measure used futures prices from the Chicago Board of Trade. Daily futures prices were obtained from the International Grain Council. A second measure is the local spot price obtained from the World Bank's commodity price database. A third measure is the world price index for each crop from the International Financial Statistical series.

Farmers make most production decisions before planting a crop. Therefore, the relevant price to include in the model is the price before planting. However, planting dates differ across countries and crops. To account for this, we use the spot price one month before the average planting date of the respective crop. Crop-specific planting dates by country were obtained from the FAO's GIEWS. However, there are some countries in our database that do not have a planting date specified. We interpolated these missing planting dates spatially using inverse distance weighting.

Prices were aggregated across crops to give an average annual grain price for each country. The price data were aggregated using a simple and weighted average across crops. The weights used to estimate the weighted average represent the share of grain production for the specific crop between 1990 and 1999. Futures prices and spot prices are measured in $/MT and the world price index (WPI) is an index with a base year of 2010.

Weather

Weather data at the country level are obtained from the Climatic Research Unit (CRU) at the University of East Anglia. CRU creates monthly country-level data by averaging monthly gridded weather data within each country. We calculate the cumulative precipitation (mm) and the average temperature (°C) for each year.

Summary statistics

The summary statistics reported in Table 1 are from the final data set used to forecast production that spans from 1990 to 2021. While the IFSA covers 77 countries, the dataset used to forecast production only includes 73 countries due to missing or erratic production data values (Cabo Verde, Congo, Jamaica, Namibia). The data are aggregated across grains. It is important to highlight that not all the countries produce the same crops so even if some crop price measures are global, they will differ across countries due to different crops produced and different planting seasons.

The average production in our dataset is 5.76 million metric tons with an average growing area of 3.30 million hectares and average yield of 1.48 metric tons per hectare. The average futures or spot prices is around $300-500 per metric ton. The world price index has a base year of 2010, so the average price index is around 200. The average precipitation is 1,135 mm/year and the average temperature is 22.18°C.

Table 1. *Summary Statistics of Data Used in Analysis (1990-2021)*

Variable	**Units**	**Mean**	**Std. Dev.**
Dependent Variables			
Grain Production	Million Metric Tons	3.30	2.92
Growing Area	Million Hectares	2.25	1.88
Grain Yield	Metric Tons per Hectare	1.81	0.95
Predictors			
Futures Price (Simple Average)	$/Metric Ton	527.94	279.01
Futures Price (Weighted Average)	$/Metric Ton	689.749	530.90
Spot Price (Simple Average)	$/ Metric Ton	369.97	157.77
Spot Price (Weighted Average)	$/ Metric Ton	8863.51	5845.36
World Price Index (Simple Average)	Index (base 2010)	217.77	93.52
World Price Index (Weighted Average)	Index (base 2010)	360.95	239.40
Precipitation	mm/year	1592.00	1097.50
Temperature	°C	26.40	5.60

.

Results

Model Selection

Using the time-series cross validation method described previously, we calculate a root mean squared error (RMSE) and mean absolute error (MAE) prediction accuracy for every possible model. The Poisson models with the best prediction accuracy for area and yield performed worse than the best linear OLS models, so we only considered linear OLS models. The best models with natural cubic spline trends also did worse than models with linear trends, so we only considered models with linear trends. There were 91 different OLS model specifications with linear trends that we considered based on different predictors included and pooling versus country-specific coefficients.

Figure 2 and Figure 3 show the RMSE and MAE across all linear models to predict growing area. A lower RMSE or MAE indicates better model performance. The models were ranked in terms of RMSE, where the model with the lowest RMSE was ranked first. The graph of MAE shows models according to the rank of RMSE. The model that had the best RMSE predictive accuracy did not necessarily have the best accuracy according to MAE. There were three models that had a much larger RMSE that were dropped from the figures to show the differences in prediction accuracy more clearly. Similarly, Figure 4 and Figure 5 show the RMSE and MAE across models to predict yield.

Two key observations stand out from these figures. First, about half of the models considered perform much worse. Prediction accuracy is roughly 25% better for the first half of the models. Second, there are minor differences in prediction accuracy among the top 5 models. While these figures cannot show what aspects of the model specification improve prediction the most, we evaluated the results and found that allowing for country-specific trends and country fixed effects are especially important for prediction accuracy.

Figure 2. *Time series cross-validation RMSE for different specifications to predict area.*

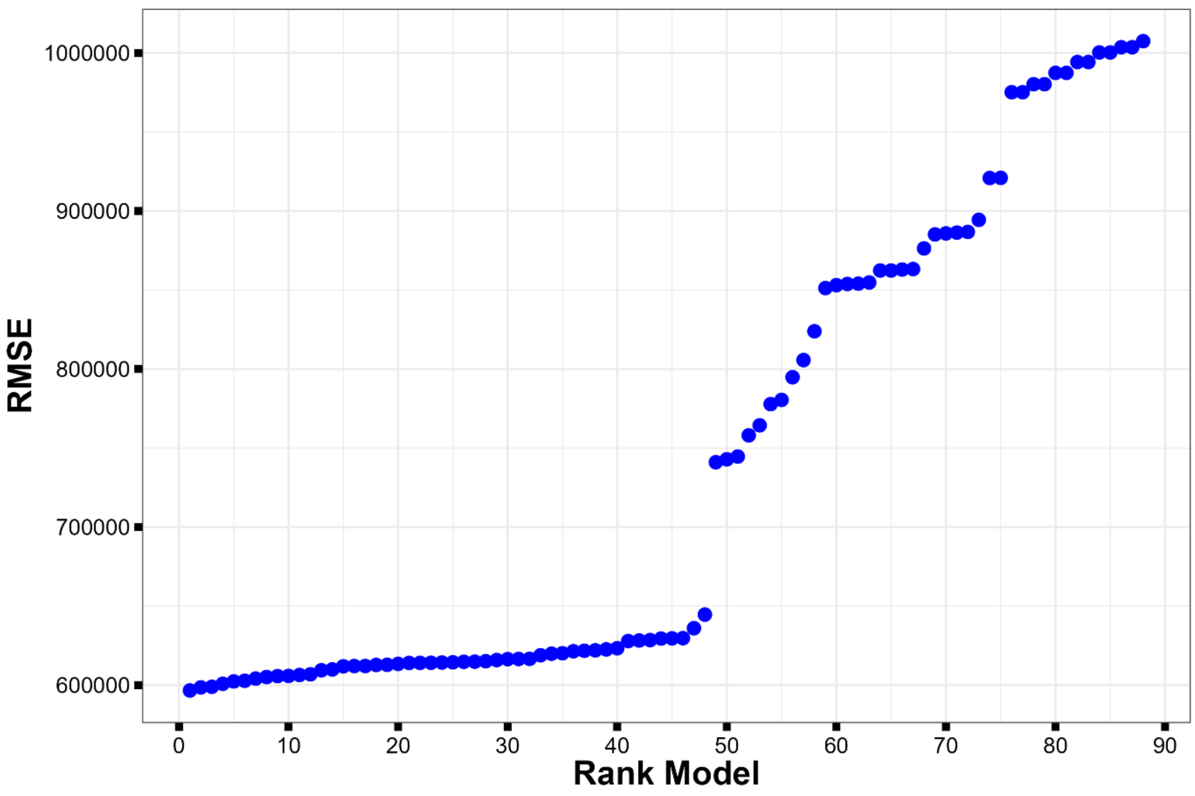

***Figure** 3. Time series cross-validation MAE for different specifications to predict area.*

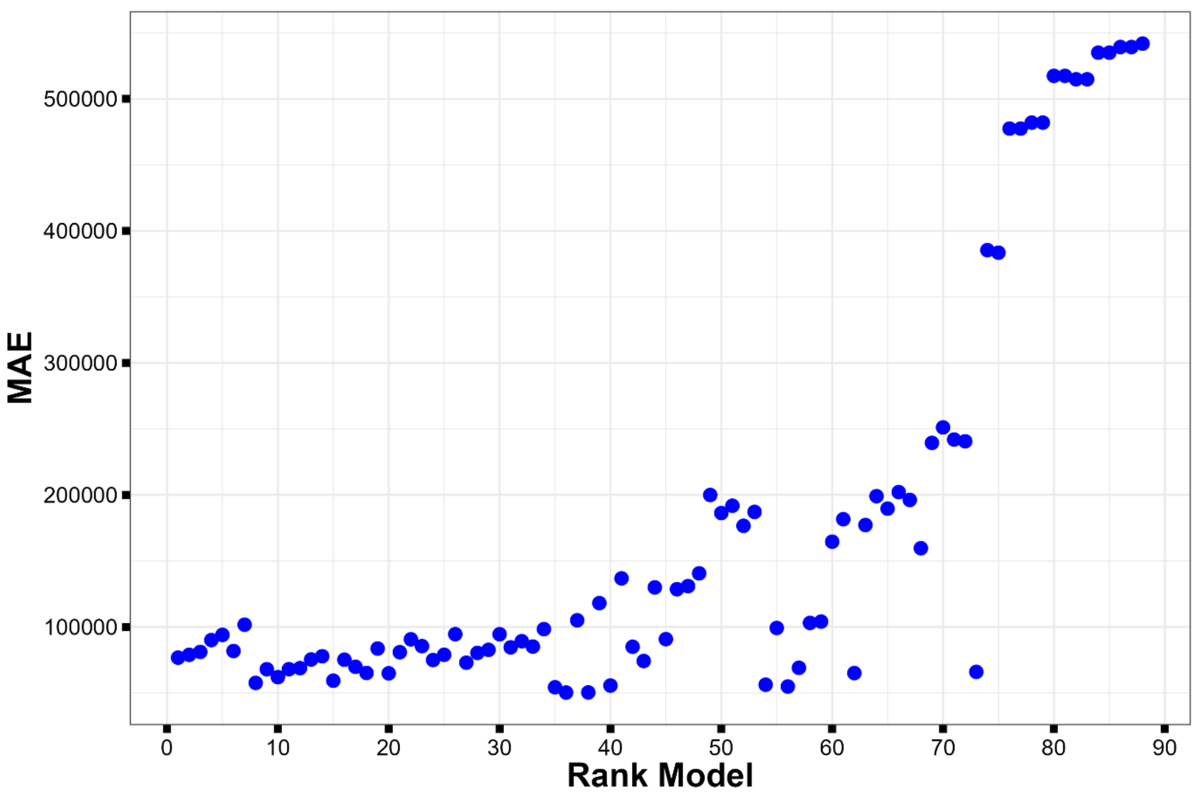

Figure 4. *Time series cross-validation RMSE for different specifications to predict yield.*

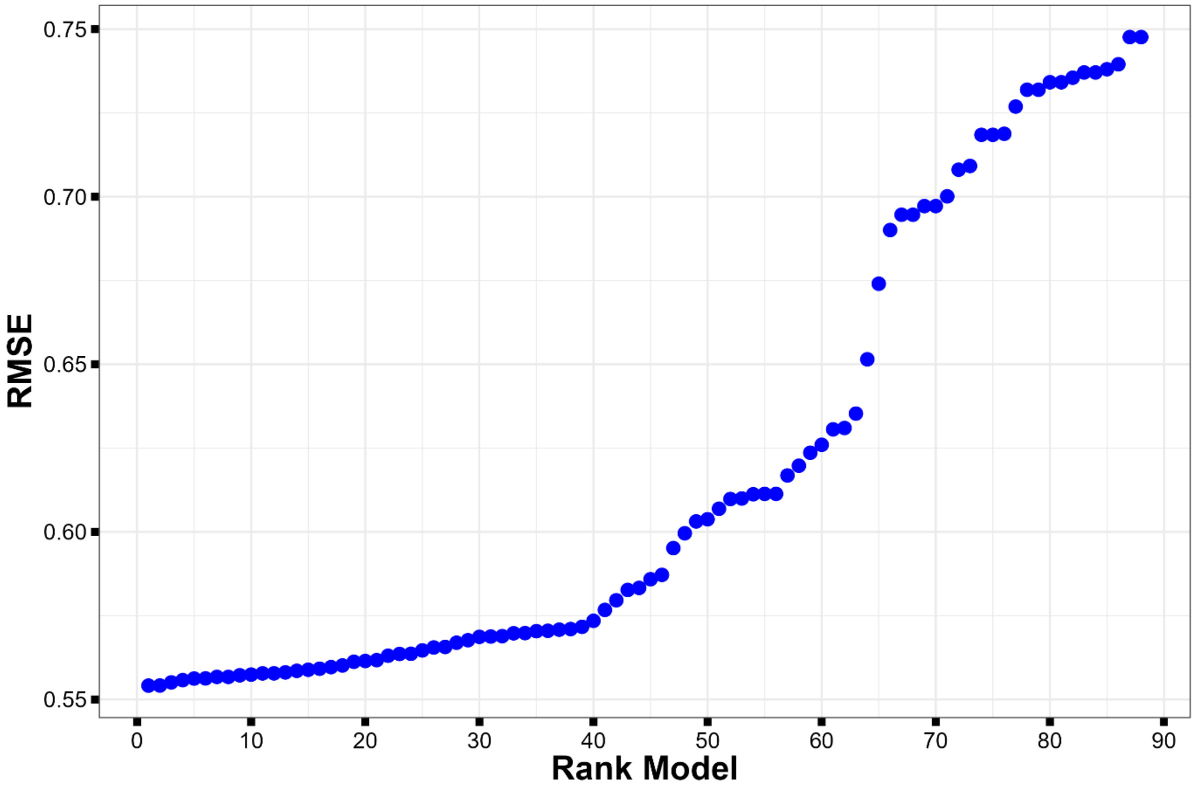

Figure 5. *Time series cross-validation MAE for different specifications to predict yield.*

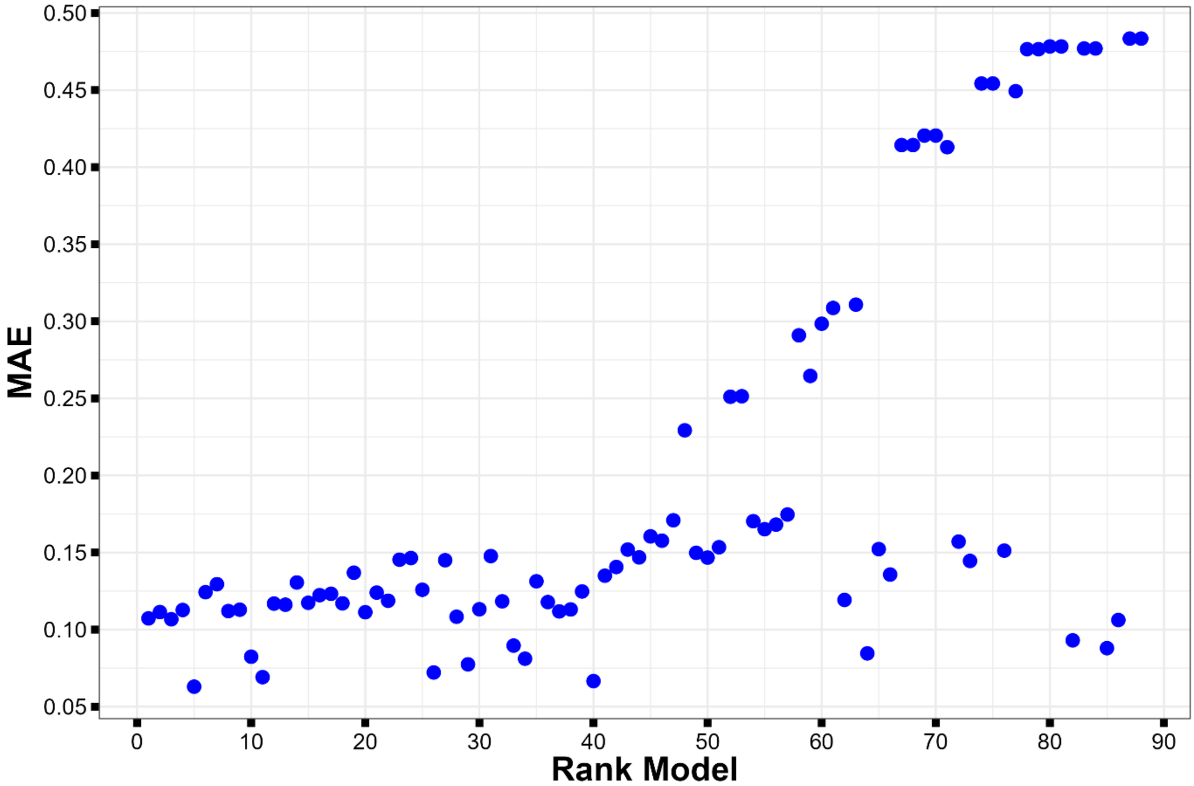

Next, we predict total production using the predictions of area and yield. Since the prediction accuracy is similar across the top 5 models, we consider all combinations of the top 5 models to predict area and the top 5 models to predict yield. This gives a total of 25 possible model combinations to predict production. However, we found that the best prediction of production was obtained by using the model that ranked first for area and first for yield. We also compared this indirect method of predicting production to a direct method of predicting production, where we predict production with each of the possible model specifications rather than area and yield. We found that the best model to predict production directly did not perform as well as the indirect method of predicting production via predictions of area and yield.

The model that performed best to predict growing area is written as

$$Area_{it} = \theta_i t + \beta_{1,i} Prec_{it} + \beta_2 FutPr_{it} + \alpha_i + \epsilon_{it}, \qquad (6)$$

where $Prec_{it}$ denotes annual precipitation and $FutPr_{it}$ denotes the futures price that is a simple average across crops. Note that the model to predict growing area includes country-specific linear trends, country-specific coefficients on precipitation, pooled coefficients on the futures price, and country-specific fixed effects. This model had a RMSE of 596,608 and a MAE of 79,747. MAE is easiest to interpret, and it indicates that on average the difference between the 10-year ahead prediction and the actual growing area was 79,747 hectares.

The model that performed best to predict yield is written as

$$Yield_{it} = \theta_i t + \beta Temp_{it} + \alpha_i + \epsilon_{it}, \qquad (7)$$

where $Temp_{it}$ is the average temperature. The model to predict yield includes country-specific linear trends, pooled coefficients on temperature, and country-specific fixed effects. This model had a RMSE of 0.554 and a MAE of 0.107. The MAE of the prediction indicates that on average the difference between the 10-year ahead prediction and the actual yield was 0.107 metric tons per hectare.

Figure 6 shows the 10-year predictions for 2012 through 2021 that are used in the time-series cross validation for our preferred model specification for every country. The vertical lines show the end of the training data sets used to create the prediction. For example, the data from 1990 to 2002 was used to predict area and yield in 2012, then these predictions were used to calculate the predicted production in 2012 as shown in green in Figure 6. Next, data from 1990 to 2003 was used to predict production in 2013.

The out-of-sample prediction accuracy is shown by comparing the green line to the blue line. The figures illustrate the challenge in forecasting production 10 years in the future. For some countries, the trend in production in the training data up to 10 years prior to the prediction was leading in a different direction than actual production after the training period. This illustrates the importance of utilizing time-series cross validation to select the preferred model specification. The model predicts especially well in countries that generally have the largest production, such as Cambodia, India, Indonesia, Malawi, Nigeria, and Pakistan. Countries that are less politically stable and where data quality may be a concern generally have poorer prediction accuracy.

Figure 6. *Times series cross-validation for the preferred specification.*

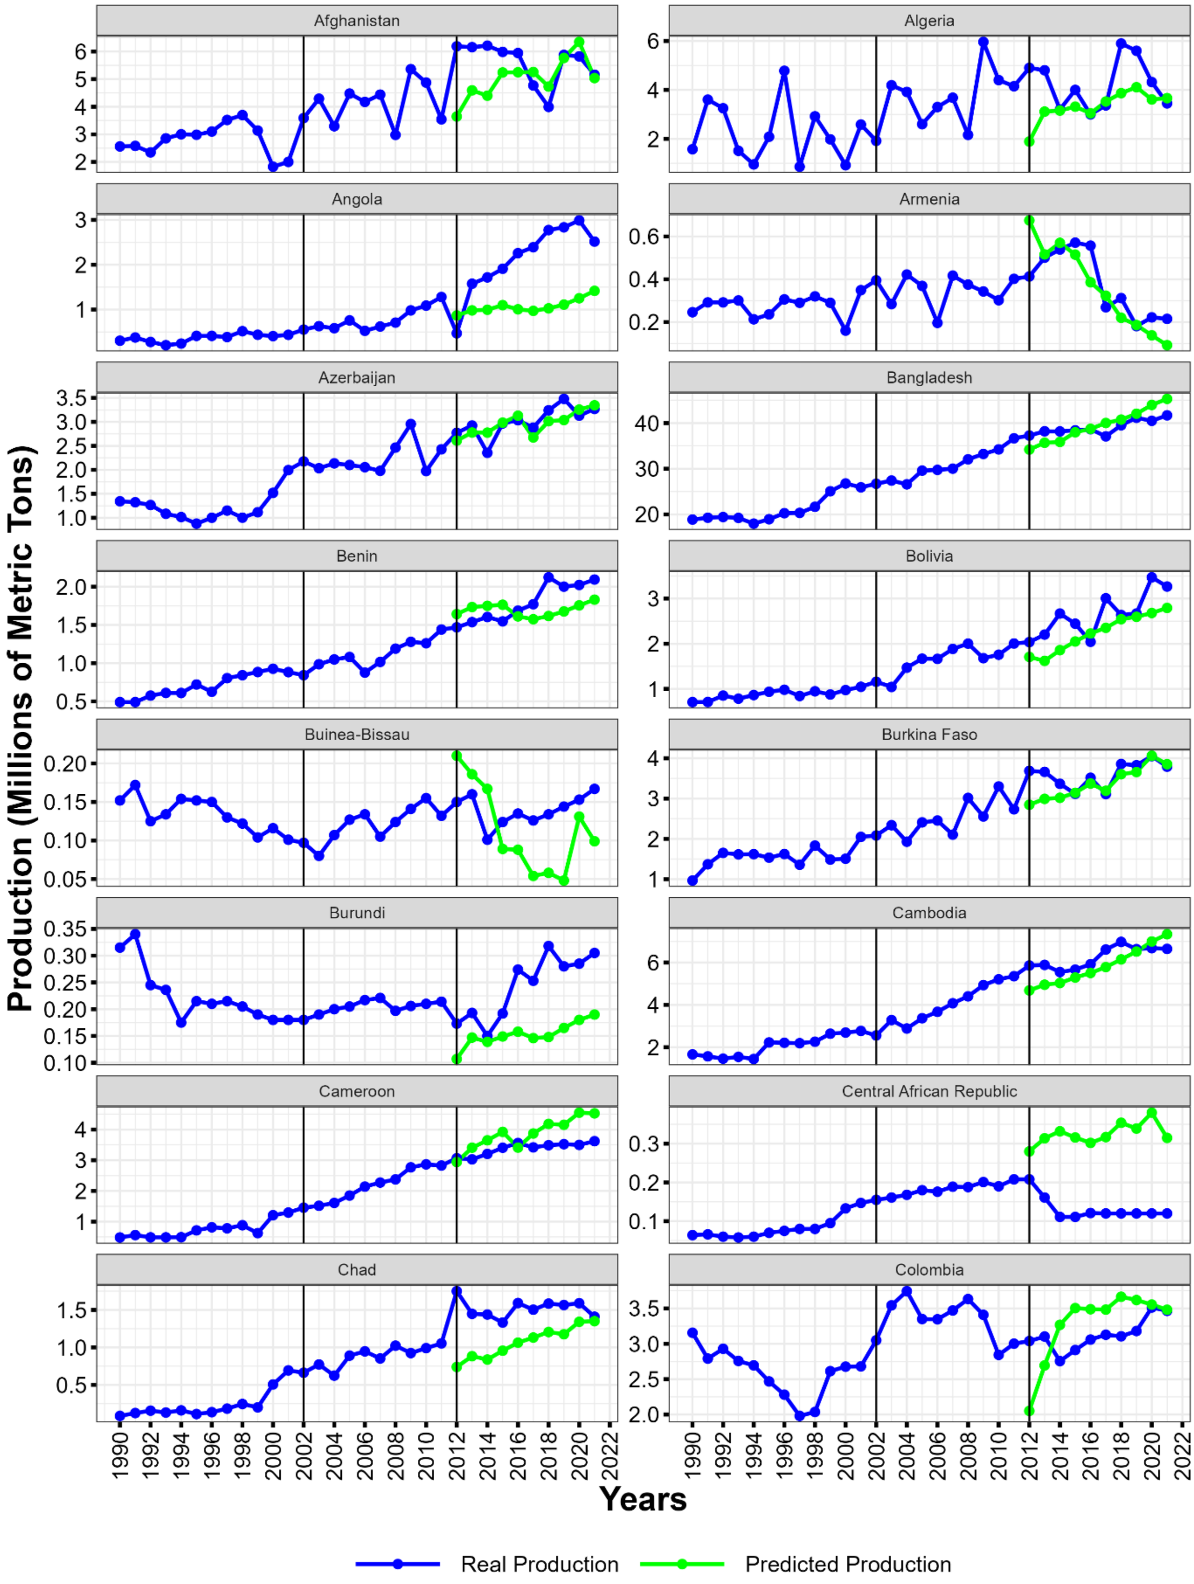

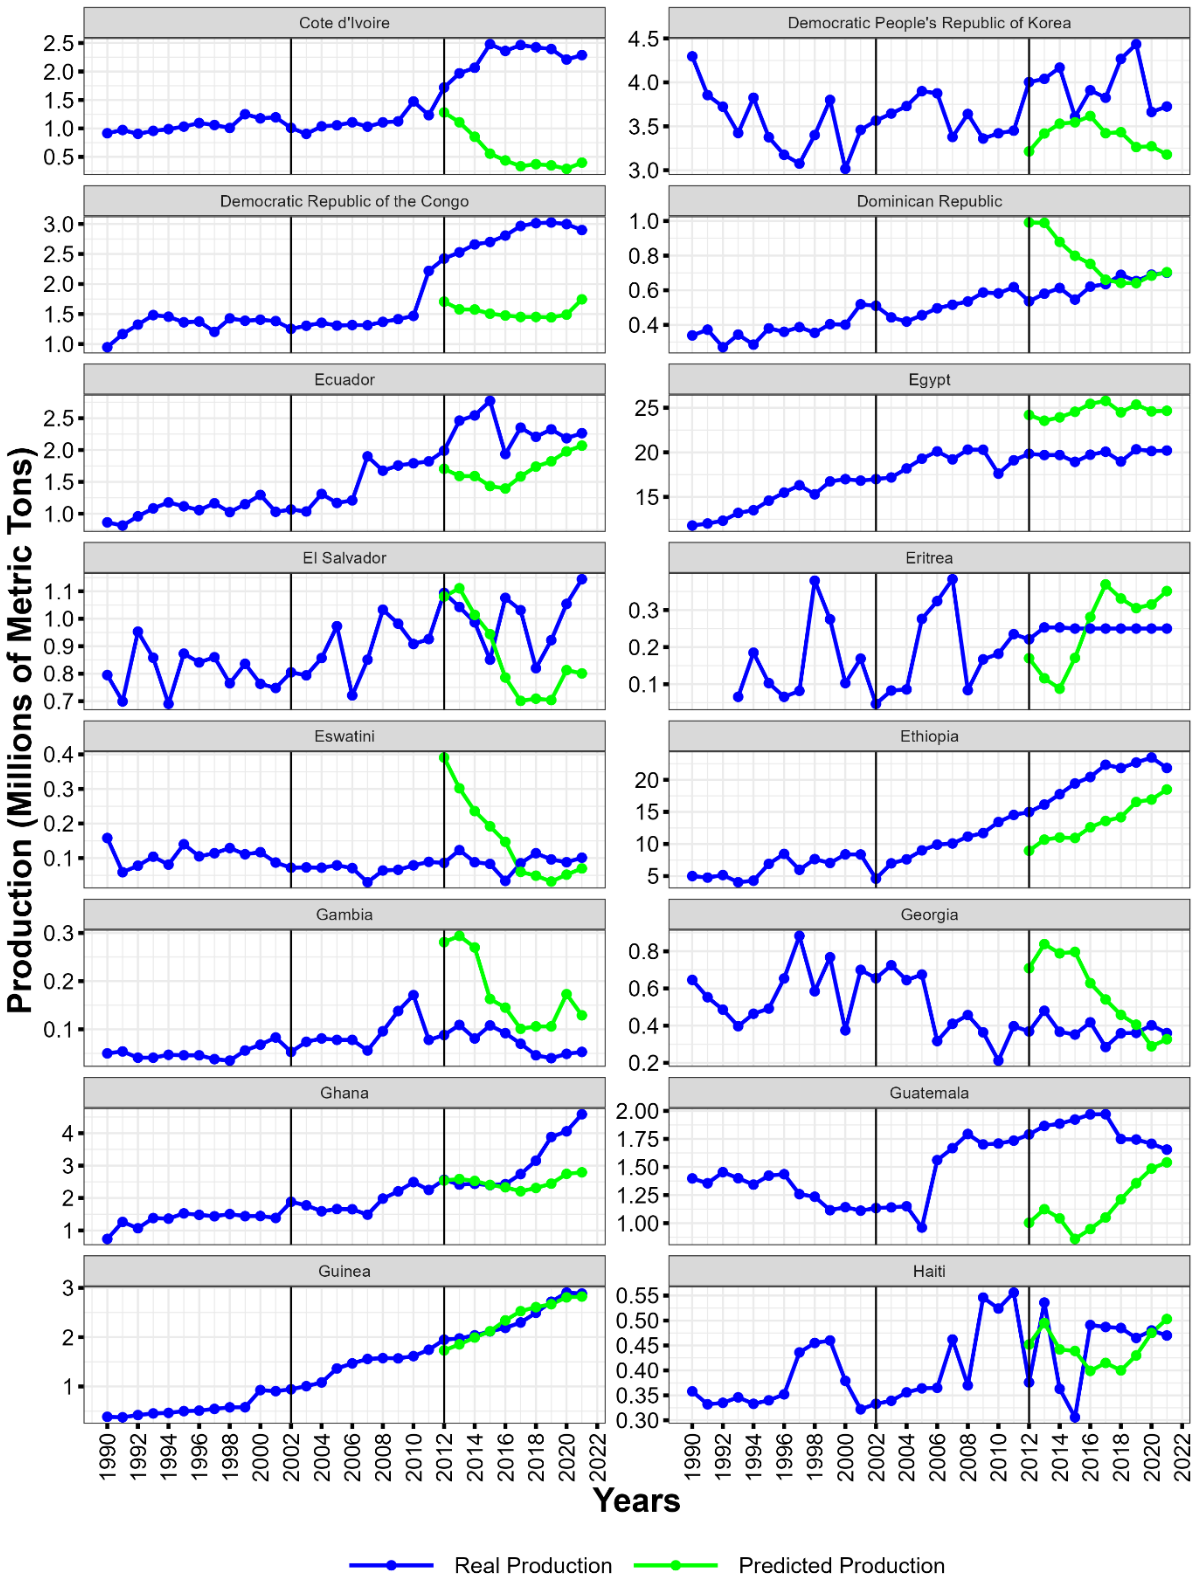

Cote d'Ivoire
Democratic People's Republic of Korea
Democratic Republic of the Congo
Dominican Republic
Ecuador
Egypt
El Salvador
Eritrea
Eswatini
Ethiopia
Gambia
Georgia
Ghana
Guatemala
Guinea
Haiti
Production (Millions of Metric Tons)
Years
Real Production
Predicted Production

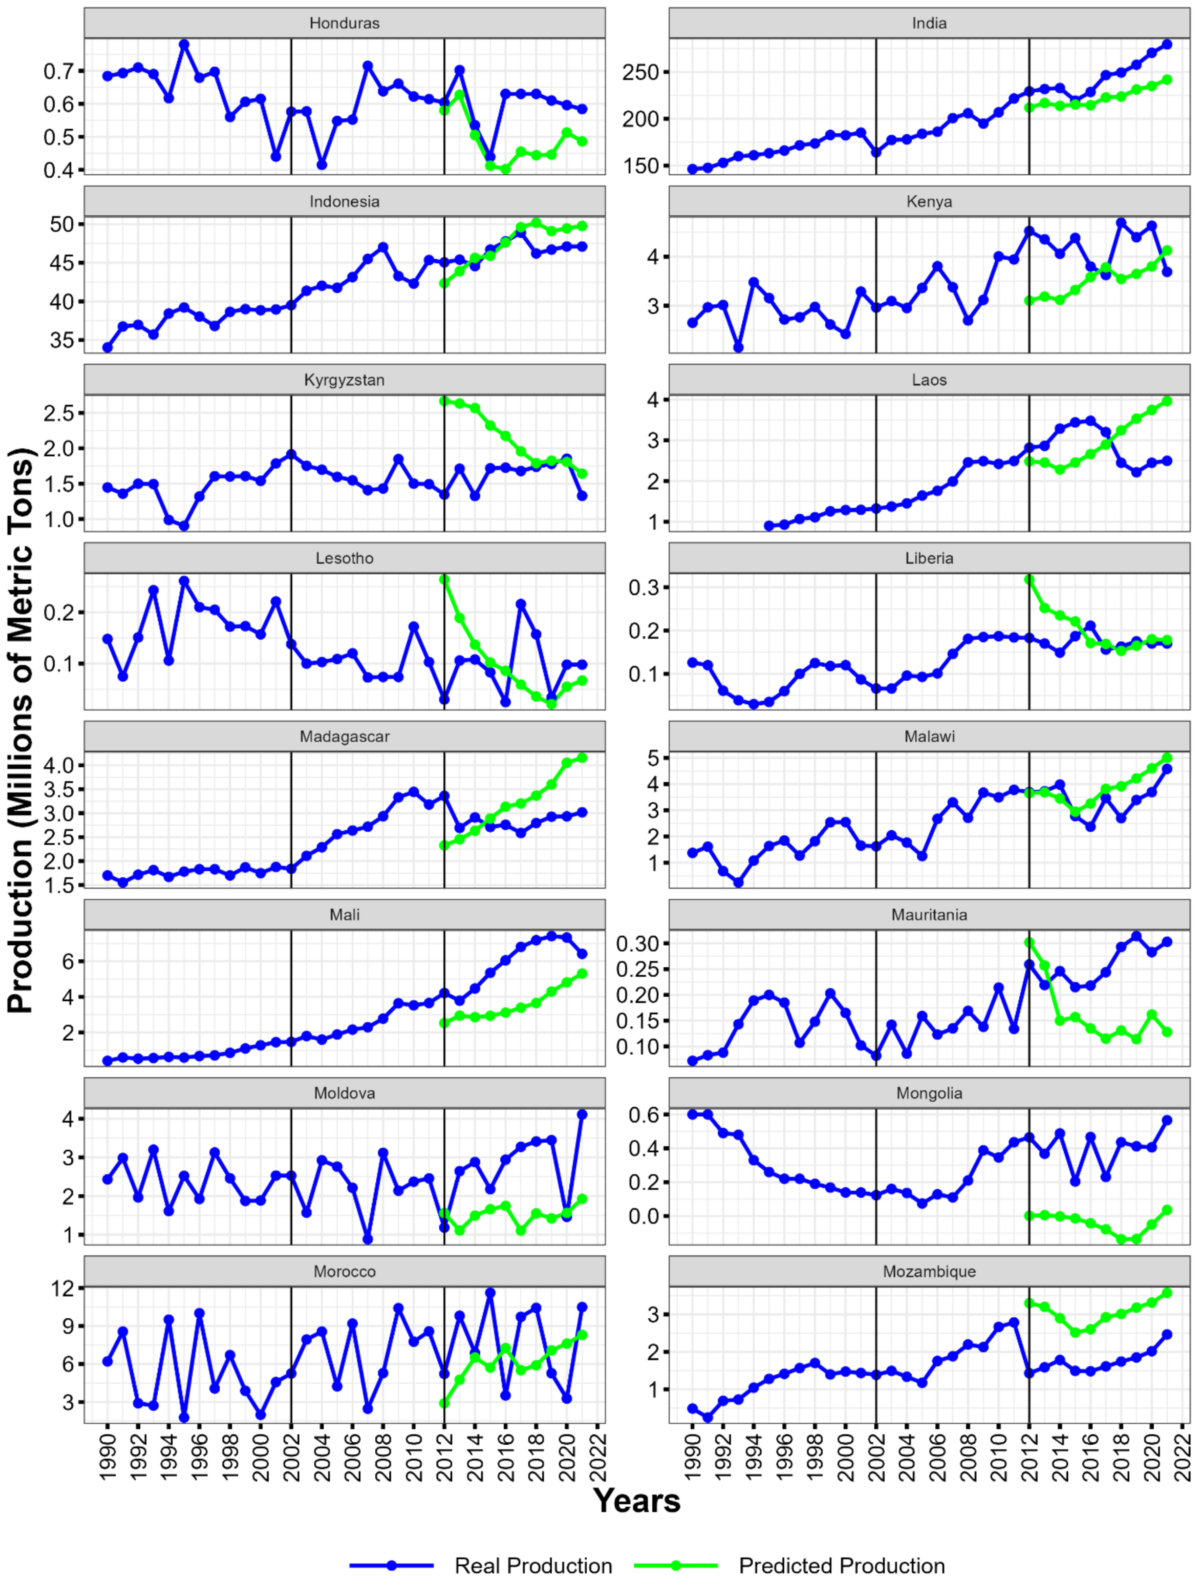

Honduras
India
Indonesia
Kenya
Kyrgyzstan
Laos
Lesotho
Liberia
Madagascar
Malawi
Mali
Mauritania
Moldova
Mongolia
Morocco
Mozambique
Production (Millions of Metric Tons)
Years
Real Production
Predicted Production

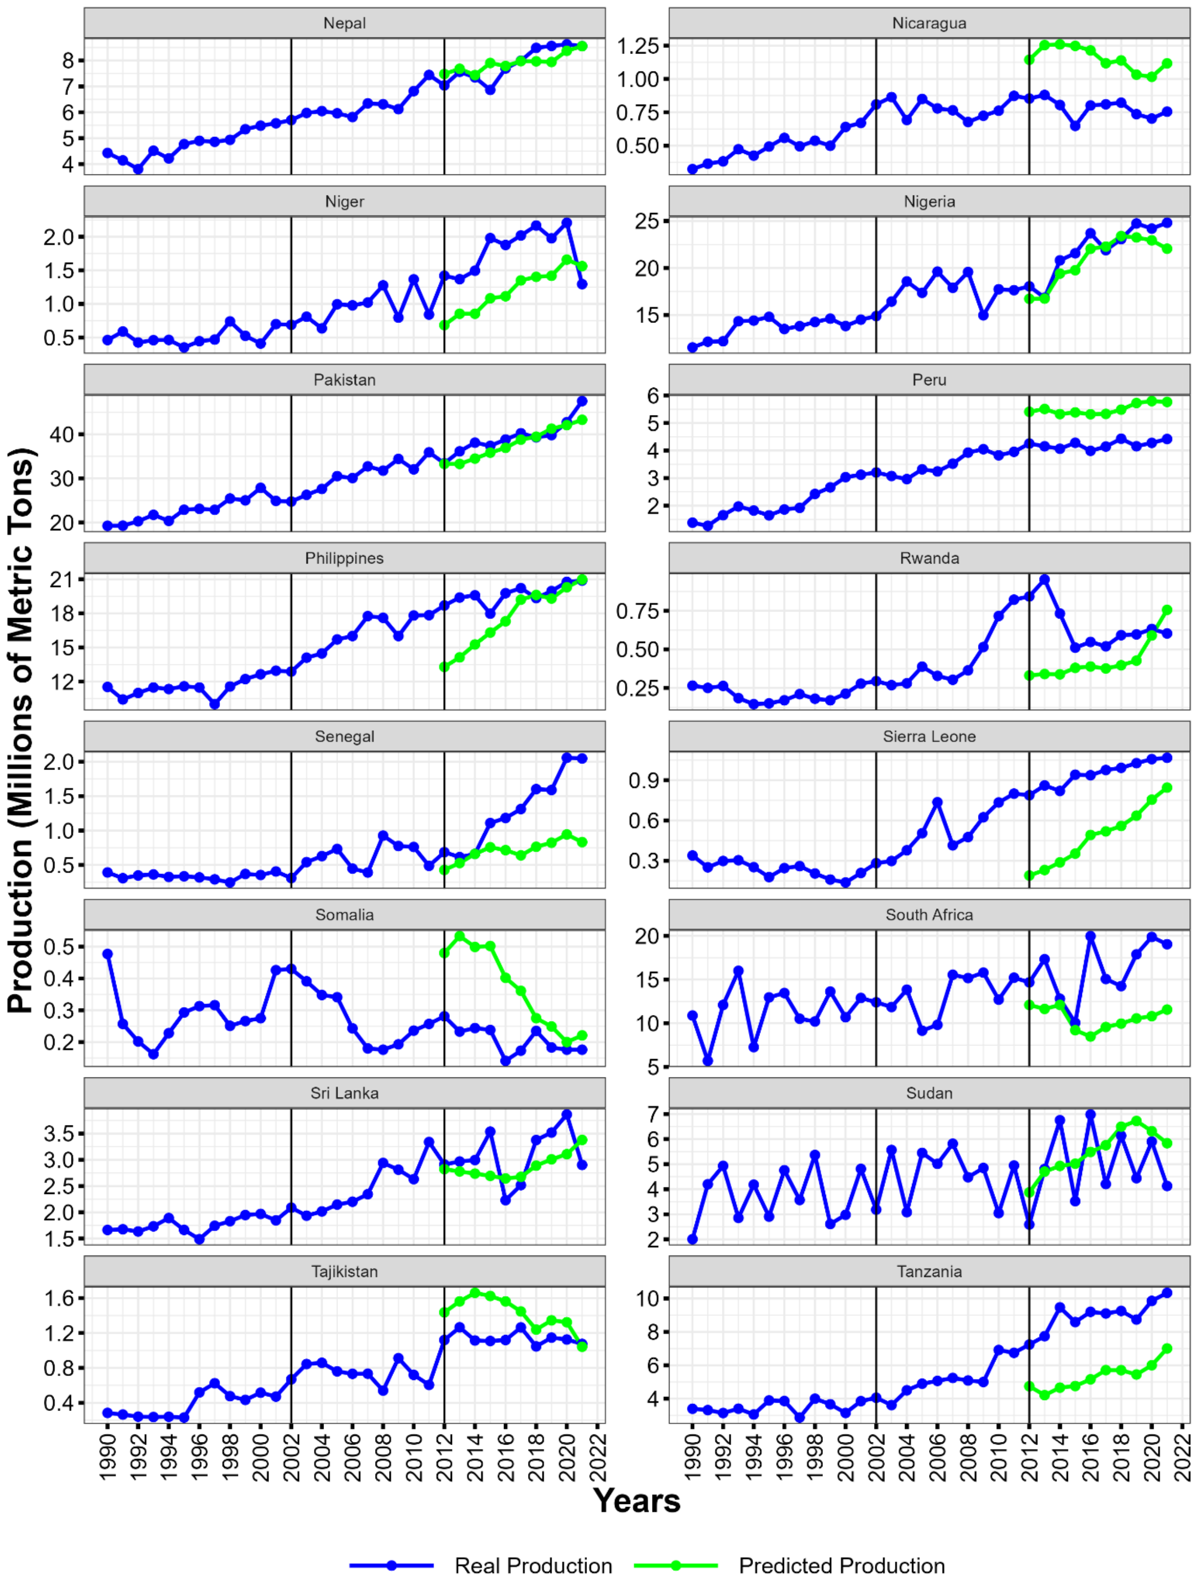

Nepal
Nicaragua
Niger
Nigeria
Pakistan
Peru
Philippines
Rwanda
Senegal
Sierra Leone
Somalia
South Africa
Sri Lanka
Sudan
Tajikistan
Tanzania
Production (Millions of Metric Tons)
Years
Real Production
Predicted Production

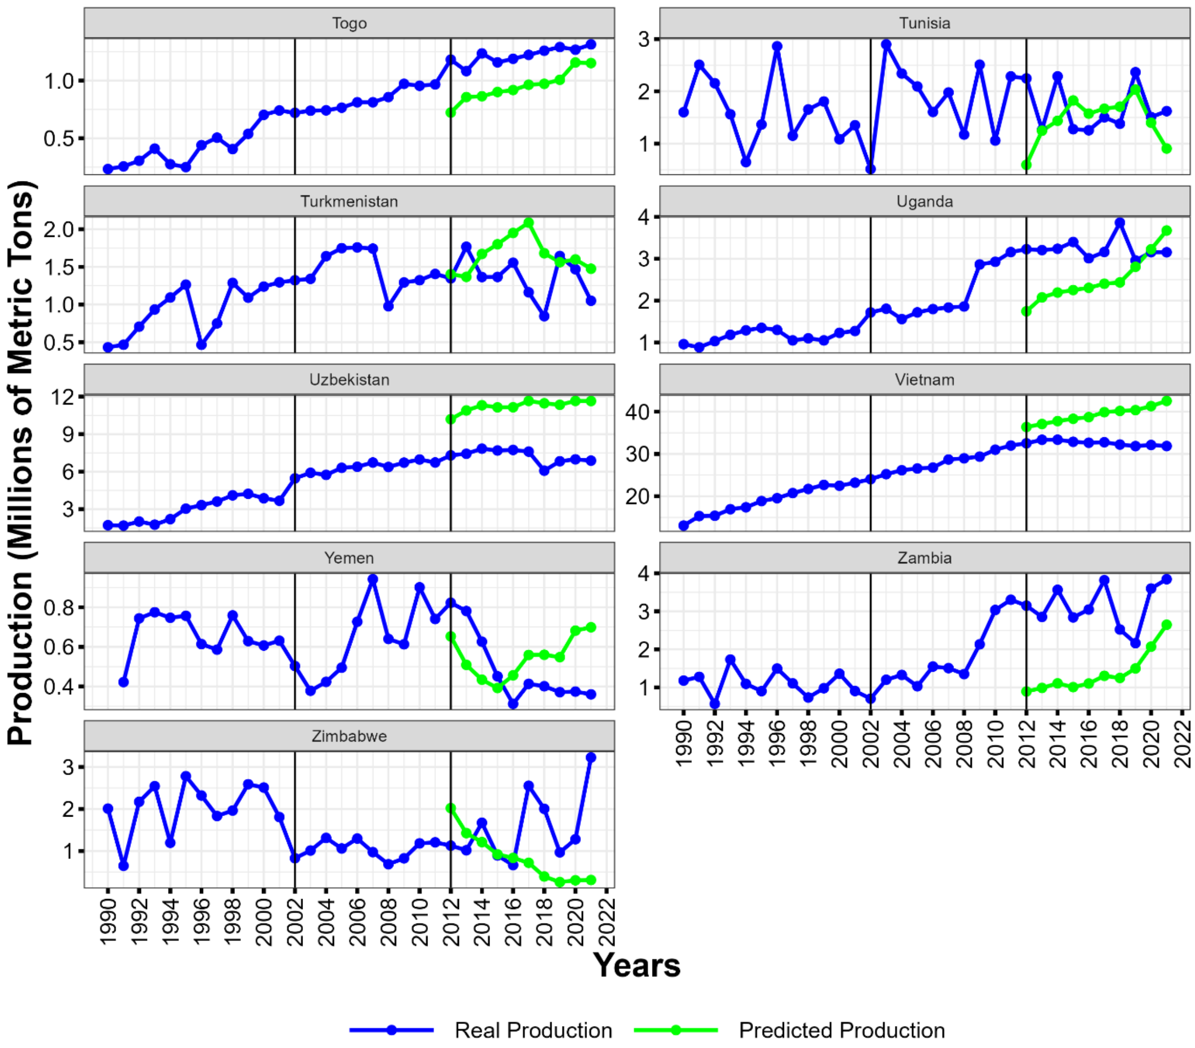

Figure 7 shows the forecasts for 2023 and 2033 production based on the preferred model specification to predict area and yield. To create these forecasts, the models were trained with all data from 1990 to 2021. Actual production in the historical data are indicated with blue and the within-sample predictions of the model are indicated with green in Figure 7. The point forecasts for 2023 and 2033 are indicated with red points. The forecasts for 2023 and 2033 assume average prices and weather over the most recent five-year observed period. Production is forecasted to increase in most countries, but there are some countries where production is forecasted to decrease over the next 10 years.

Figure 7. Production Forecasts for 2023 and 2033.

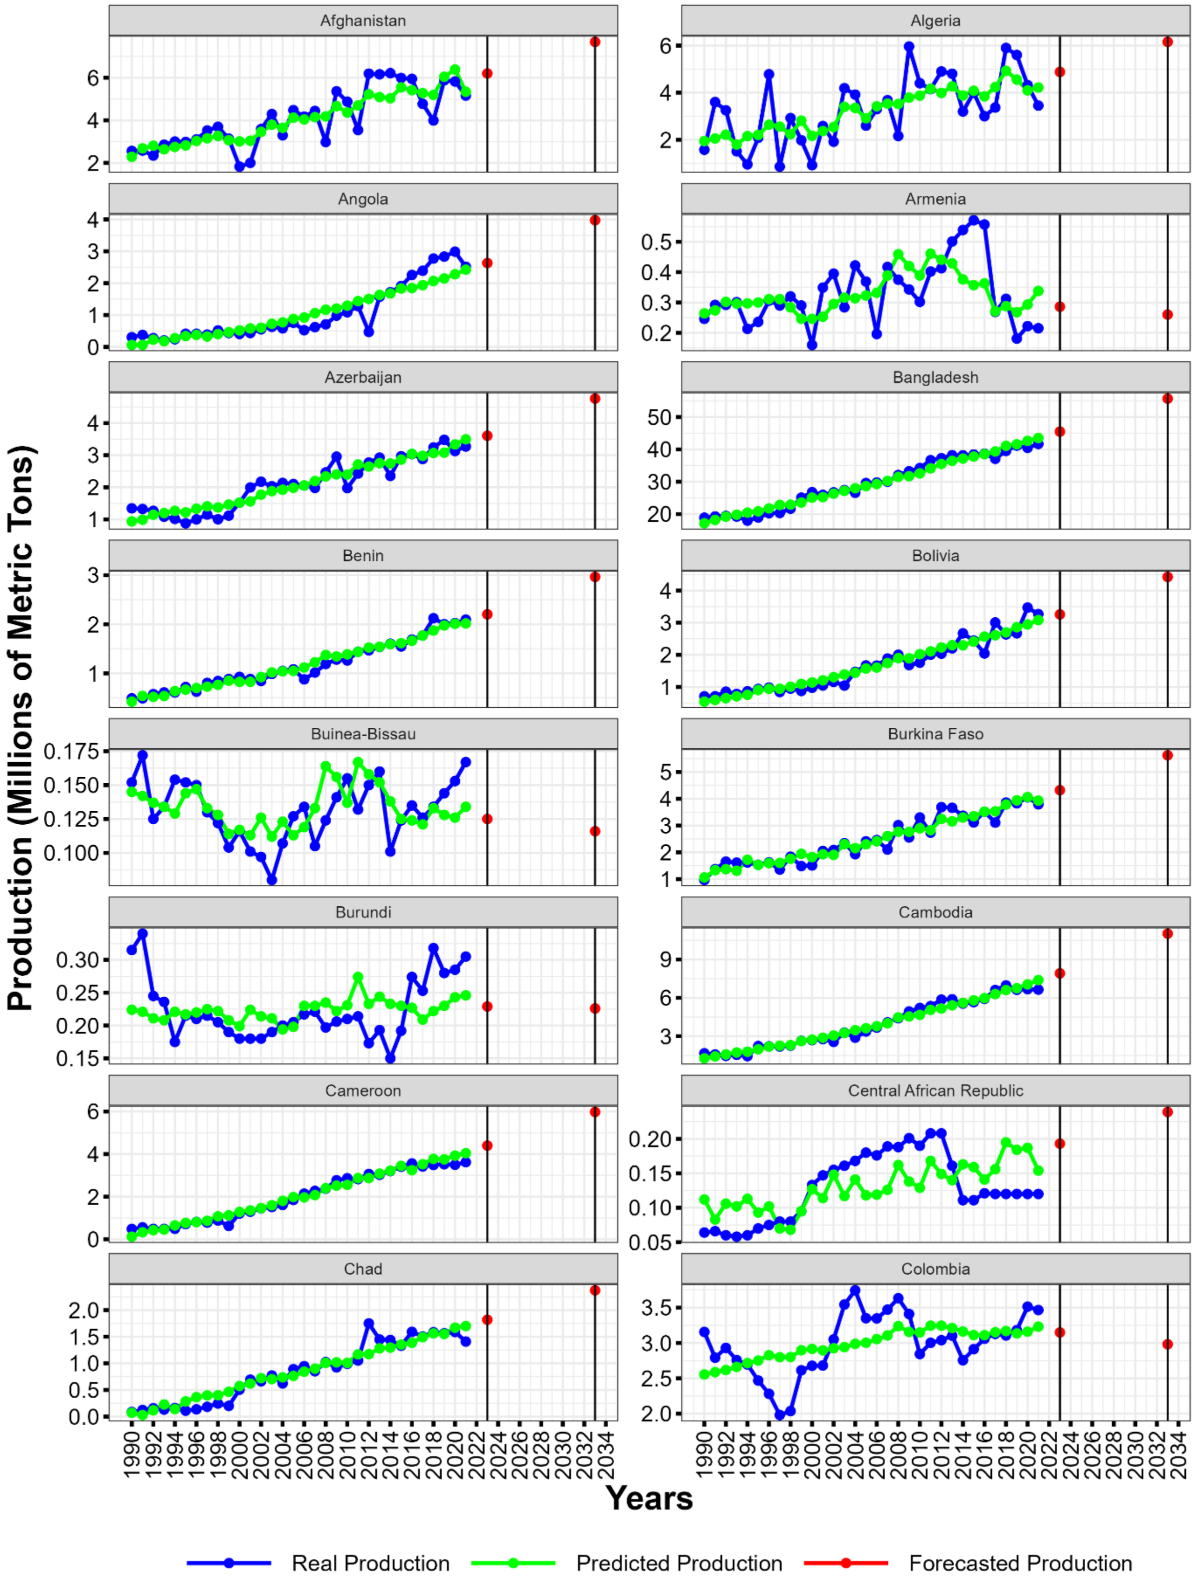

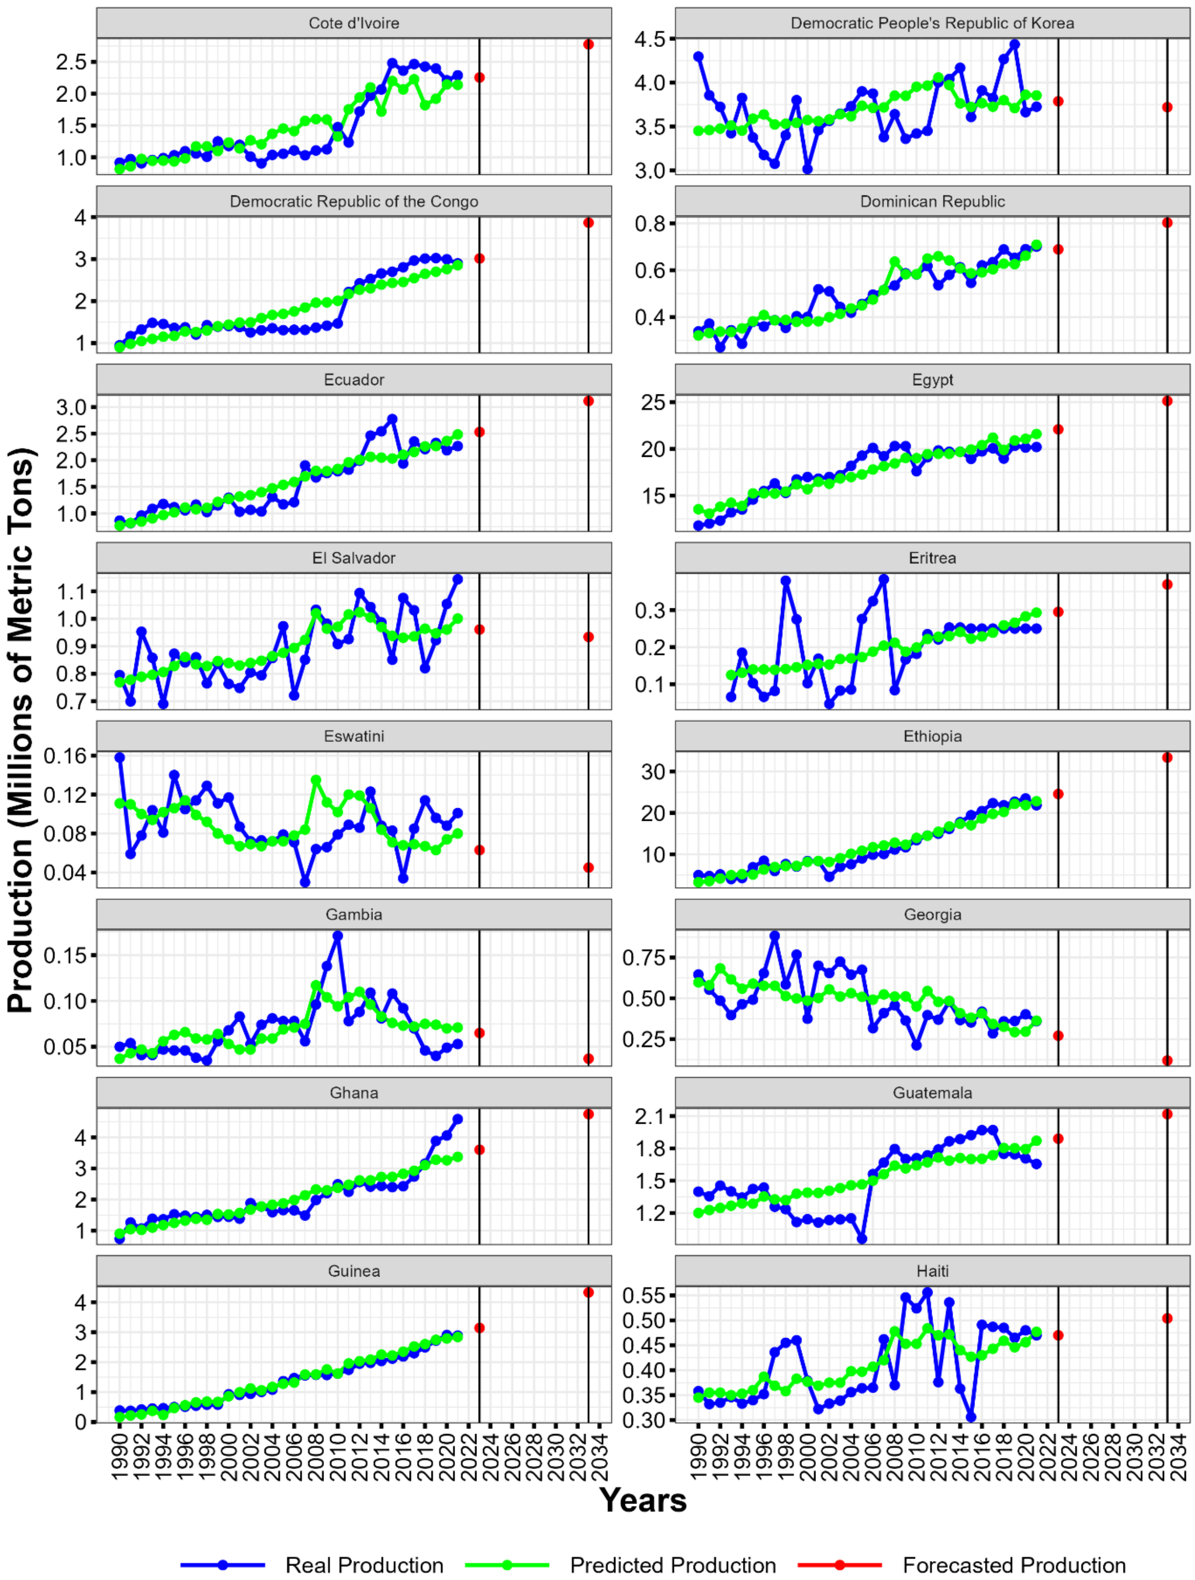

Cote d'Ivoire
Democratic People's Republic of Korea
Democratic Republic of the Congo
Dominican Republic
Ecuador
Egypt
El Salvador
Eritrea
Eswatini
Ethiopia
Gambia
Georgia
Ghana
Guatemala
Guinea
Haiti
Production (Millions of Metric Tons)
Years
Real Production
Predicted Production
Forecasted Production

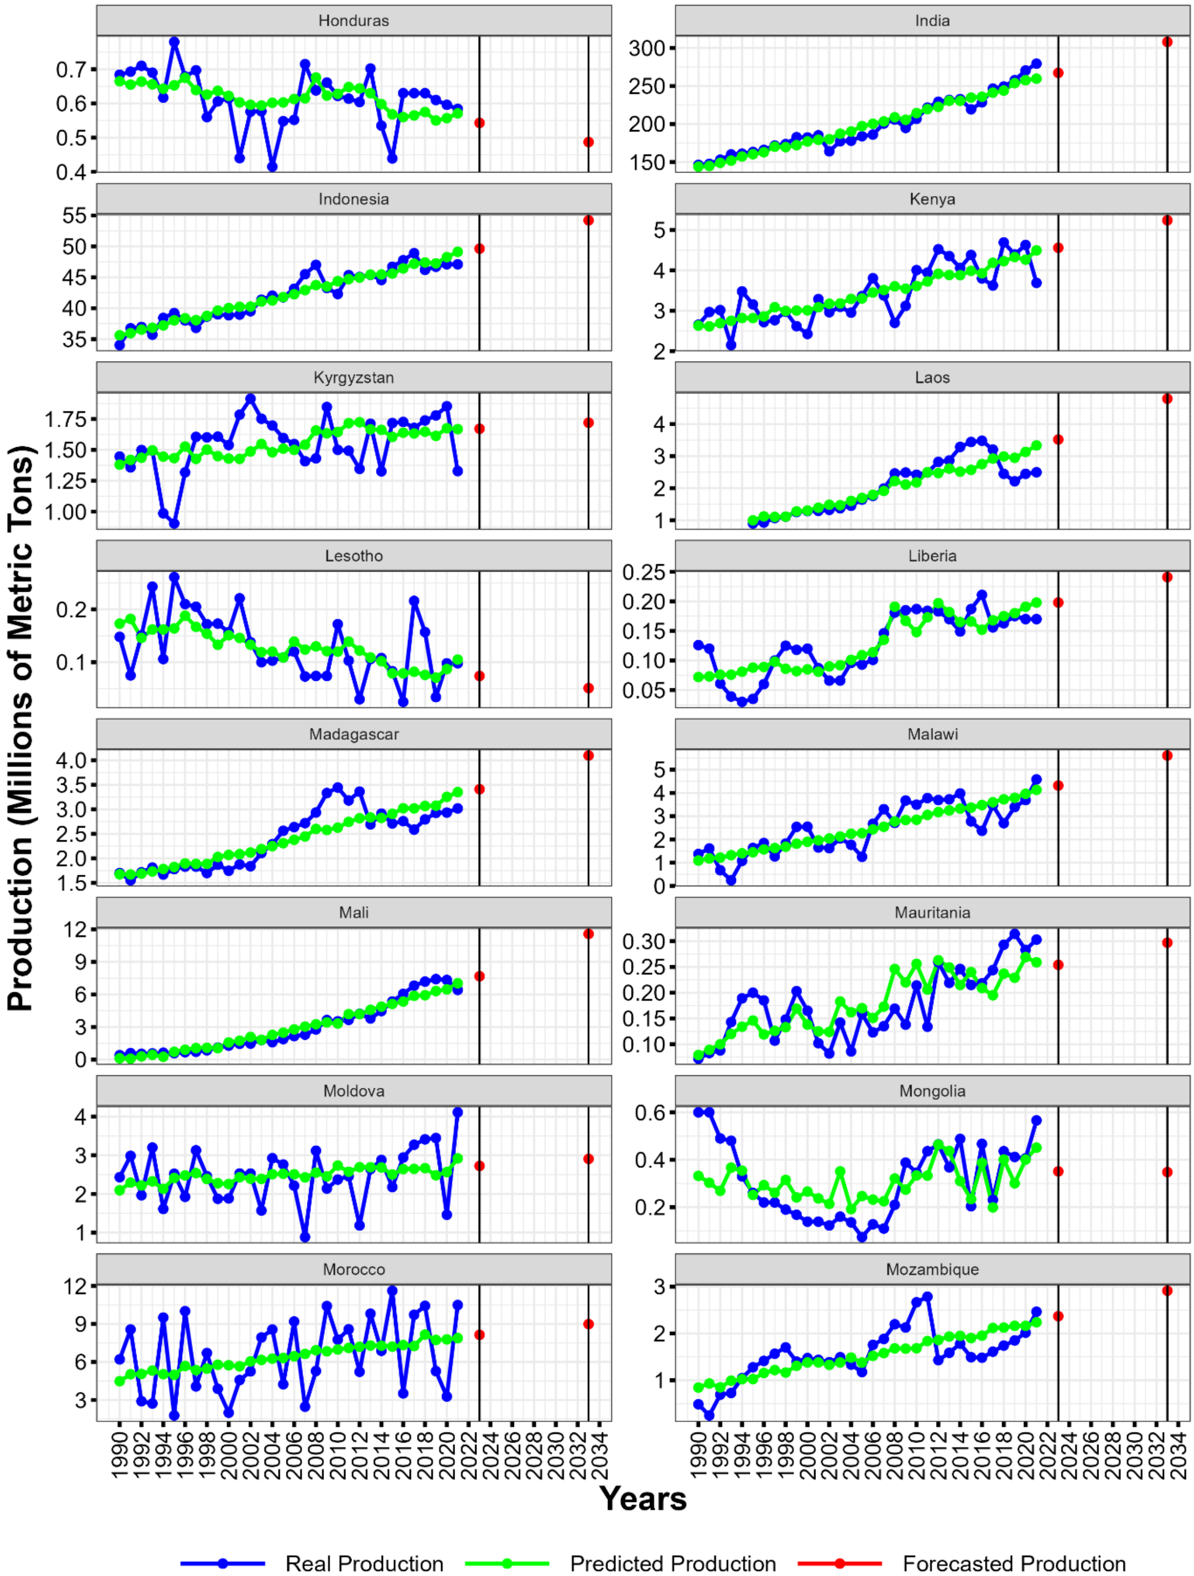

Honduras
India
Indonesia
Kenya
Kyrgyzstan
Laos
Lesotho
Liberia
Madagascar
Malawi
Mali
Mauritania
Moldova
Mongolia
Morocco
Mozambique
Production (Millions of Metric Tons)
Years
Real Production
Predicted Production
Forecasted Production

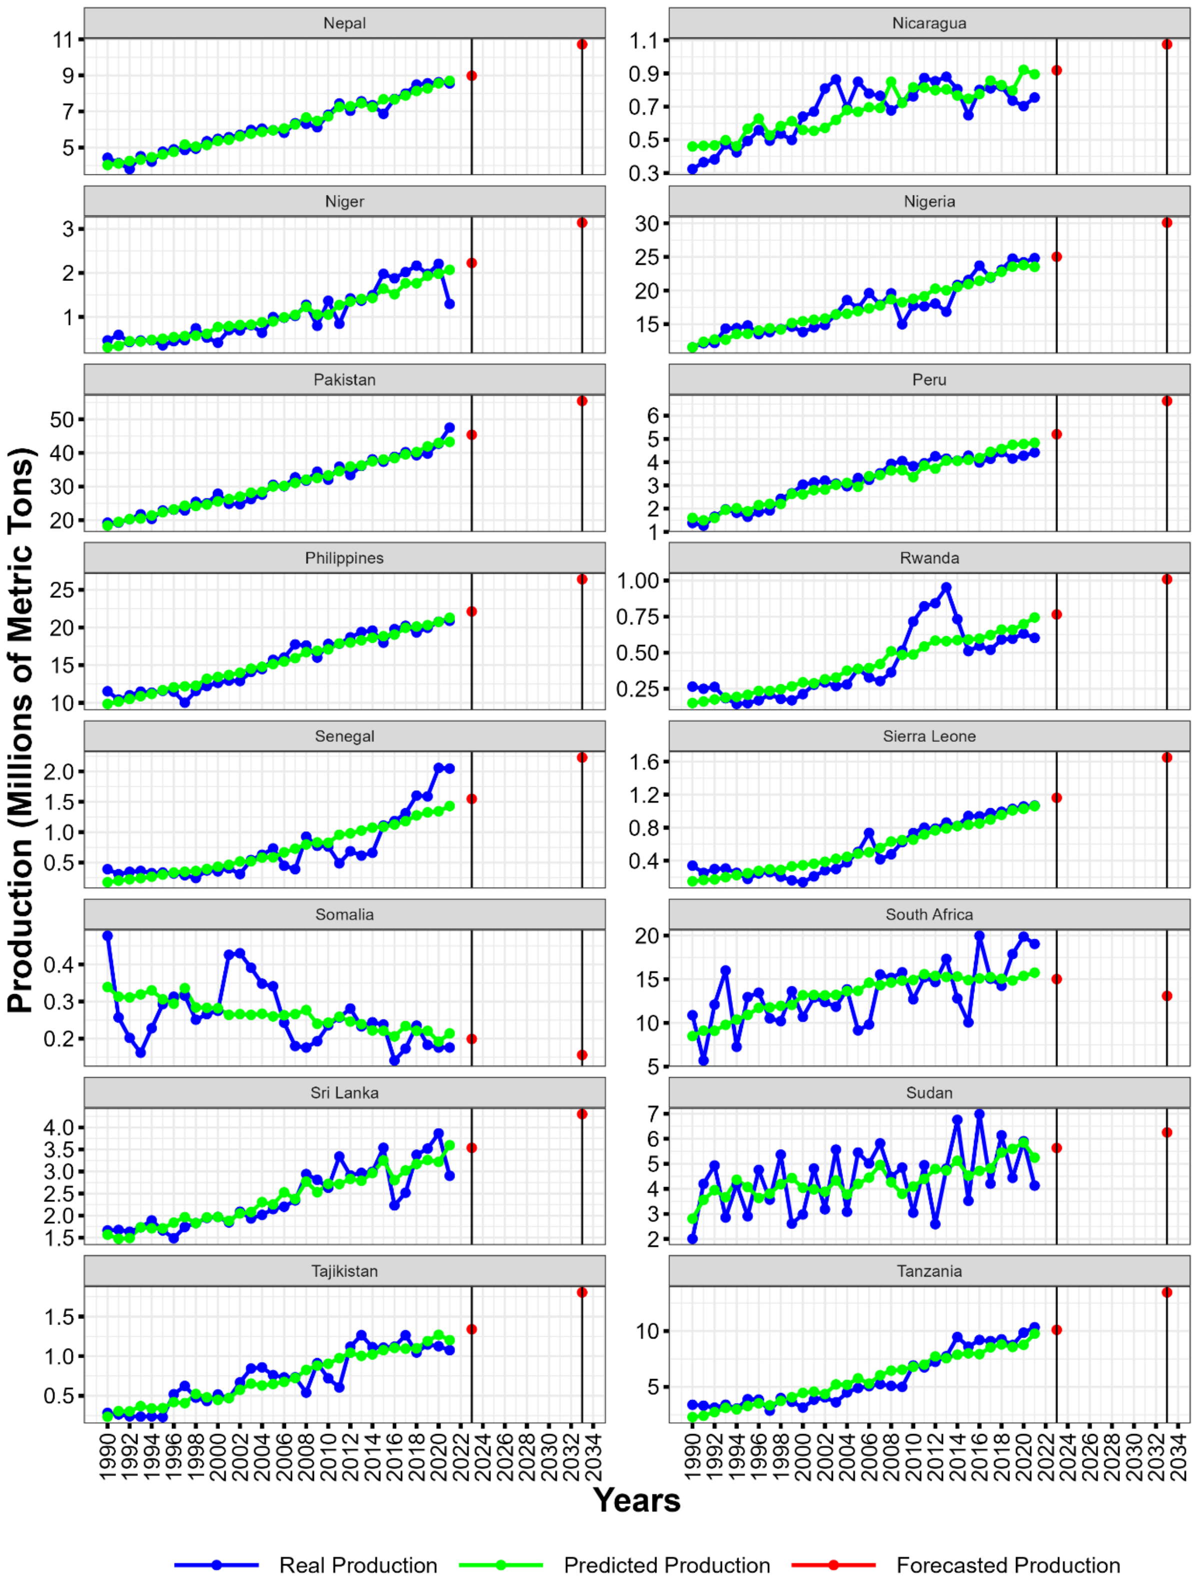

Nepal
Nicaragua
Niger
Nigeria
Pakistan
Peru
Philippines
Rwanda
Senegal
Sierra Leone
Somalia
South Africa
Sri Lanka
Sudan
Tajikistan
Tanzania
Production (Millions of Metric Tons)
Years
Real Production
Predicted Production
Forecasted Production

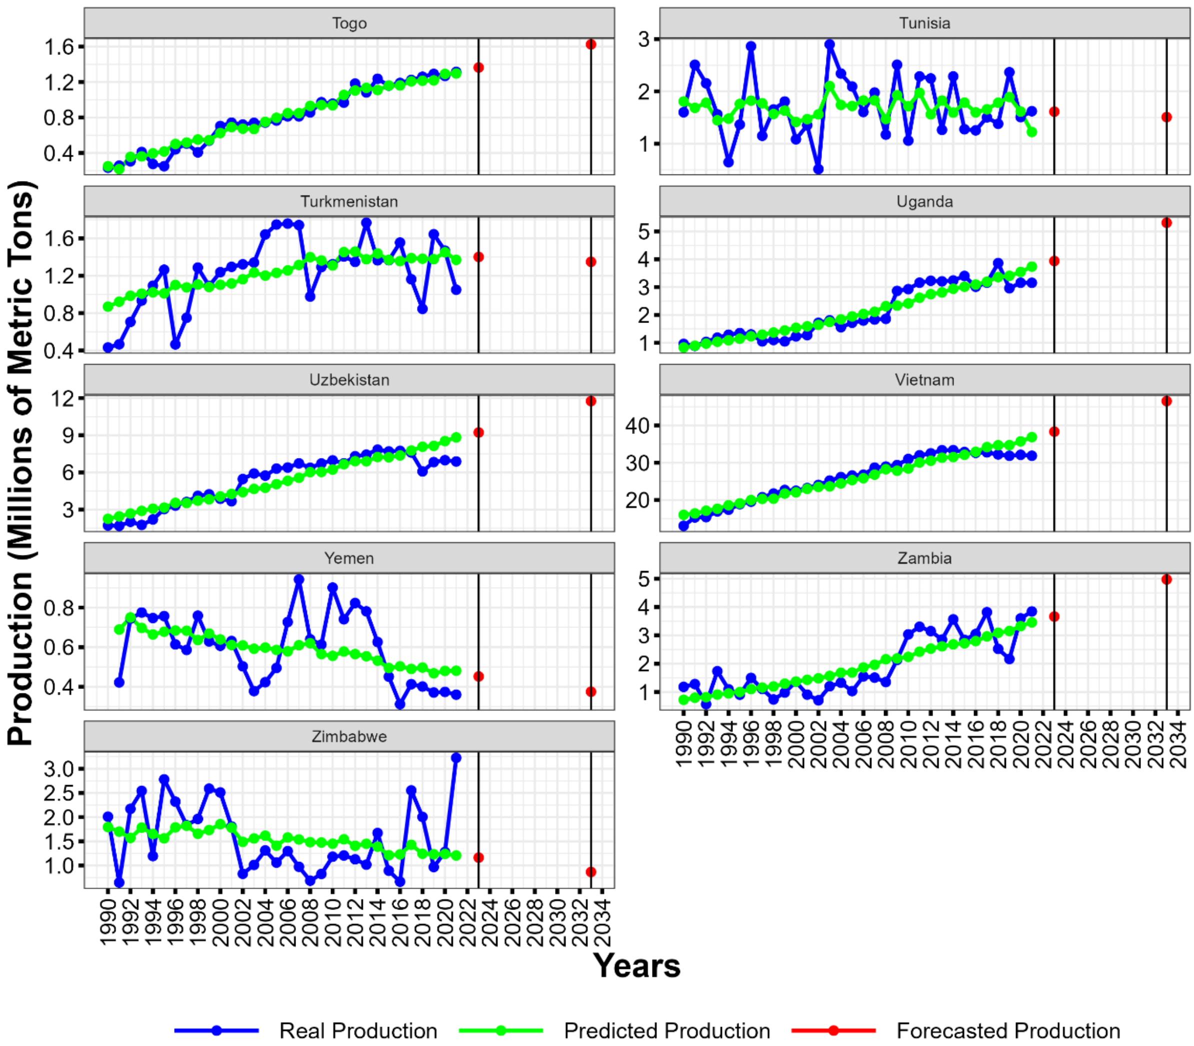

The previous IFSA report used an alternative method to forecast production. Next, we briefly explain the previous method and compare the forecast accuracy to the revised method. The previous IFSA model did not forecast growing area with an econometric model. Instead, an area equation was calibrated to the prior 3 years of data. Yield was forecasted using the following equation:

$$Yield_{it} = \theta_r t + \beta_{1,r} MA_{2,it} + \beta_{2,r} MA_{5,it} + \alpha_i + \epsilon_{it}, \qquad (8)$$

where the subscript r denotes 10 different subregions of countries, $MA_{2,it}$ denotes a moving average of yield over the previous 2 years, and $MA_{5,it}$ denotes a 5-year moving average. Key differences with the revised specification is that the previous model included subregion-specific trends rather than country-specific trends, the previous model included moving averages, and the previous model did not include weather predictors.

To compare the revised and previous methods to forecast yield, we conducted time-series cross validation for a 10-year ahead forecast with both models. Since the previous model included a 5-year moving average, the training data could only begin in 1995 rather than 1990. Using this timeframe for both methods, we estimate accuracy metrics from the revised method as RMSE: 0.525 and MAE: 0.0653. Accuracy from the previous method were RMSE: 0.540 and MAE: 0.0647. The revised method performed better than the previous method in terms of RMSE and the two methods were similar in terms of MAE for forecasting yield. However, another key advantage of the revised methodology is that it forecasts growing area.

Conclusion

Our study improves the current IFSA model used to forecast production 1 and 10 years into the future. The accuracy measures indicate that the indirect method performs best at forecasting production. The best model specification to predict and forecast growing area includes country-specific linear trends, annual precipitation, and the futures price that is a simple average across crops, and country-specific fixed effects. While the best model specification to forecast yield includes country-specific linear trends, pooled coefficients on temperature, and country-specific fixed effects.

When comparing our best model specification with the previous methods from the IFSA model to predict yield, the results indicate that our revised models performed better than the previous method in terms of RMSE, and the two methods were similar in terms of MAE for forecasting yield. However, another key advantage of the revised methodology is that it forecasts growing areas.

References

Baquedano, F. G., Zereyesus, Y. A., Valdes, C., & Ajewole, K. (2021). International Food Security Assessment 2021-31. *USDA Economic Research Service, 1962*.

Grassini, P., Eskridge, K. M., & Cassman, K. G. (2013). Distinguishing between Yield Advances and Yield Plateaus in Historical Crop Production Trends. *Nature Communications, 4(1), 2918.*

Haile, M. G., Kalkuhl, M., & von Braun, J. (2016). Worldwide Acreage and Yield Response to International Price Change and Volatility: A Dynamic Panel Data Analysis for Wheat, Rice, Corn, and Soybeans. *American Journal of Agricultural Economics., 98(1),* , 172-190. https://doi.org/10.1093/ajae/aav013.

Hendricks, N. P., Janzen, J. P., & Smith, A. (2014). Futures Prices in Supply Analysis: Are Instrumental Variables Necessary? *American Journal of Agricultural Economics, 97(1),* , 22-39. http://www.jstor.org/stable/24476999

Hyndman, R. J., & Athanasopoulos, G. (2021). Forecasting: principles and practice. *OTexts, 3rd Edition*(Monash university, Australia), 442. https://otexts.com/fpp3/

Lobell, D. B., Schlenker, W., & Costa-Roberts, J. (2011). Climate trends and global crop production since 1980. *Science, 333(6042)*, 616-620.

Ray, D. K., Ramankutty, N., Mueller, N. D., West, P. C., & Foley, J. A. (2012). Recent Patterns of Crop Yield Growth and Stagnation. *Nature Communications., 3(1)*, 1293.

Roberts, M. J., & Schlenker, W. (2013). Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate. *American Economic Review., 103(6)*, 2265-2295.

Zereyesus, Y. A., Baquedano, F., & Morgan, S. (2022). Evaluating the Prediction Performance of the International Food Security Assessment's Production Models: A Cross-Validation Approach. *USDA Economic Research Service, Technical bulletin Number 1959*, 24. https://www.ers.usda.gov/webdocs/publications/105126/tb-1959.pdf?v=6667.8