

# Climate-Resilience and Cereal-Yield Joint Forecasting for Low- and Middle-Income Countries with Public Data and Grounded LLM Policy Explanation

Lumeng Han<sup>1</sup>, Anna Chen<sup>2</sup>

<sup>1</sup>International Agricultural Development, University of California, Davis, CA, USA

<sup>2</sup>Industrial Engineering and Operations Research, UC Berkeley, Berkeley, CA, USA

[lmhan@ucdavis.edu](mailto:lmhan@ucdavis.edu)

DOI: 10.69987/JACS.2024.40109

## Keywords

climate resilience; cereal yield; World Development Indicators; ND-GAIN; QoG Environmental Indicators; panel forecasting; low- and middle-income countries; explainable AI; large language models

## Abstract

This paper reports a country-year forecasting experiment that links climate-resilience indicators, observed climate exposure, and cereal yields for low- and middle-income countries. The empirical panel combines the December 16, 2024 World Development Indicators archive for cereal yield, cereal production, cereal land and GDP, an ND-GAIN country-year file for vulnerability, readiness, ND-GAIN index and HDI, and QoG Environmental Indicators variables for annual temperature and rainfall. Country-specific temperature and rainfall anomalies are calculated against each country's 1981-2000 climate baseline. After complete-case filtering, the panel contains 39 countries, 702 complete rows, feature years 2002-2019, and target years 2003-2020. All features at year  $t$  are used to predict national cereal yield at year  $t+1$ . The evaluation uses a temporal split with feature years 2002-2016 for training and 2017-2019 for held-out testing, plus rolling-origin validation. Nine forecasting specifications are compared: persistence, country mean, panel fixed-effects ordinary least squares, ridge, Huber-ridge, regression tree, random forest, gradient boosting, and residual gradient boosting. The ridge panel model is the best held-out specification with RMSE 318.23 kg/ha, MAE 192.93 kg/ha, MAPE 9.08%, and  $R^2$  0.945. Recent cereal-yield history remains the strongest accuracy driver, while the climate and ND-GAIN variables add auditable exposure and adaptation context. A constrained policy-explanation layer then converts measured vulnerability, readiness, HDI, climate anomalies and yield volatility into country-specific adaptation explanations.

## Introduction

Cereal production is a central food-security indicator because wheat, rice, maize and other dry-grain cereals provide a large share of calories and rural income in many low- and middle-income countries. Climate hazards affect this system through heat stress, drought, flood damage, pest pressure, and disrupted input and market systems. Broad assessment reports and empirical crop studies show that climate exposure and adaptive capacity shape food-system risk, while temperature and

precipitation shocks can reduce national and regional yields [1], [7]-[11]. Evidence from yield-trend, crop-model, disaster and regional-impact studies also shows that aggregate yield responses vary by crop, geography and production system [12]-[18].

This study addresses the combined prediction and explanation problem with public country-year data. The data design follows three constraints. First, the panel must cover a common annual period for many low- and middle-income countries. Second, the

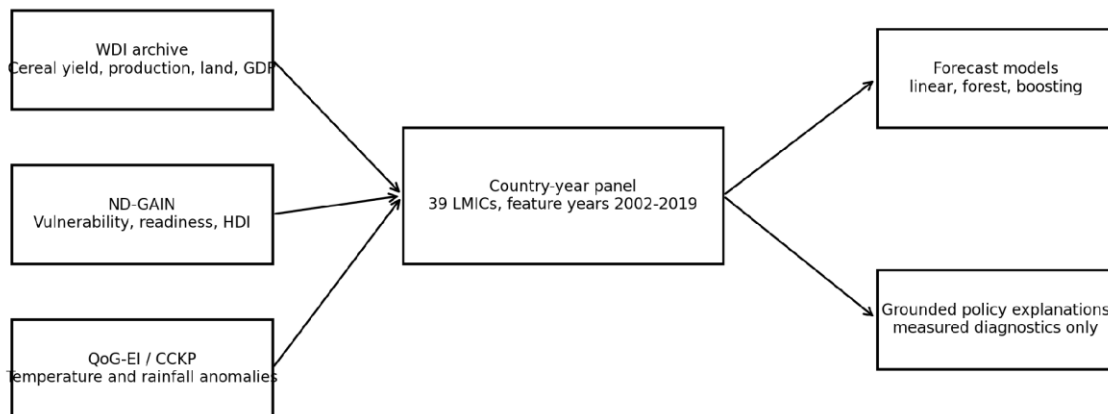
target and features must be auditable from public data releases: FAOSTAT-derived WDI cereal-yield and cereal-production indicators, WDI GDP, ND-GAIN vulnerability and readiness measures, and annual temperature and rainfall indicators distributed through the QoG Environmental Indicators file [2]-[6]. Third, the analysis must separate forecast accuracy from policy explanation. A high-accuracy one-year forecast can be driven mainly by yield persistence, but policy users also need measured evidence on vulnerability, readiness and climate exposure.

The paper makes four contributions. It builds a balanced panel of 39 LMICs for feature years 2002-2019 and target years 2003-2020. It evaluates nine forecasting specifications with a time-based test design and rolling-origin validation. It reports the roles of yield history, production/land indicators,

GDP, ND-GAIN resilience variables, and temperature and rainfall anomalies in the same panel. It then uses a constrained language-model policy interpreter that is grounded in measured country diagnostics and adaptation categories.

The empirical target is national cereal yield in kg/ha. The target is an aggregate dry-grain cereal indicator rather than separate crop-specific targets because the harmonized WDI cereal-yield, production and land indicators are consistently available for the selected country-year panel. The aggregate includes major cereal groups relevant to LMIC food systems, especially wheat, rice and maize, along with other cereals. The aggregate target therefore supports a coherent national food-security forecasting panel that can be linked to macroeconomic, resilience and climate-exposure indicators.

Joint forecasting and policy-explanation workflow



**Figure 1.** Data and modeling workflow linking WDI, ND-GAIN and QoG-EI climate indicators to forecasting and grounded policy explanation.

## Method

Figure 1 summarizes the workflow. The experiment uses three public-data blocks available by 2024. The WDI December 16, 2024 archive supplies cereal yield, cereal production, land under cereal

production and GDP. The ND-GAIN country-year file supplies vulnerability, readiness, ND-GAIN index and HDI. The QoG Environmental Indicators time-series file supplies CCKP annual temperature and rainfall variables, which are derived from climate data sources used for country-level climate indicators [3]-[6]. Table 1 lists the data blocks and variables.

**Table 1.** Data sources and variables used in the empirical panel.

Source	Raw file	Indicator code	Role	Years
WDI December 16, 2024 archive	WDIEXCEL.xlsx	AG.YLD.CREL.KG	target and lagged cereal yield (kg per hectare)	2000-2021 used
WDI December 16, 2024 archive	WDIEXCEL.xlsx	AG.PRD.CREL.MT; AG.LND.CREL.HA	cereal production and cereal land features	2002-2019
WDI December 16, 2024 archive	WDIEXCEL.xlsx	NY.GDP.MKTP.CD	GDP control, log-transformed	2002-2019
ND-GAIN country-year file	ndgain-readiness-vulnerability-allyears.txt	vulnerability; readiness; hdi	climate vulnerability, readiness and adaptation context	2002-2019
QoG Environmental Indicators time series	qog_ei_ts_sept21.csv	cckp_temp; cckp_rain	annual temperature and rainfall; anomaly features	1981-2000 baseline; 2002-2019 features

The country filter contains 39 low- and middle-income countries from Asia, Africa, Latin America, the Middle East and North Africa. A row is retained only when the cereal yield target, lagged yield variables, cereal production, cereal land, GDP, ND-

GAIN variables, HDI, annual temperature and annual rainfall are observed for the required year. The final panel has 702 complete rows. Feature years run from 2002 to 2019; target years run from 2003 to 2020. Table 2 reports country coverage after filtering.


**Table 2.** Country coverage after complete-case filtering.

ISO3	Country	Years	Obs	Mean yield	Vuln	Ready	HDI
AFG	Afghanistan	2002-2019	18	1,888.0	0.597	0.251	0.443
BEN	Benin	2002-2019	18	1,306.5	0.563	0.307	0.493
BFA	Burkina Faso	2002-2019	18	1,080.8	0.570	0.286	0.377
BGD	Bangladesh	2002-2019	18	4,200.6	0.556	0.354	0.562
BOL	Bolivia	2002-2019	18	2,008.3	0.463	0.283	0.667
BTN	Bhutan	2002-2019	18	2,734.9	0.529	0.415	0.605

CIV	Cote d'Ivoire	2002-2019	18	1,995.2	0.496	0.233	0.486
CMR	Cameroon	2002-2019	18	1,764.2	0.470	0.260	0.520
EGY	Egypt, Arab Rep.	2002-2019	18	7,198.1	0.426	0.327	0.679
ETH	Ethiopia	2002-2019	18	1,878.5	0.559	0.292	0.410
GHA	Ghana	2002-2019	18	1,675.1	0.482	0.344	0.574
GTM	Guatemala	2002-2019	18	2,125.6	0.454	0.291	0.607
HND	Honduras	2002-2019	18	1,678.0	0.463	0.273	0.596
IDN	Indonesia	2002-2019	18	4,798.2	0.454	0.397	0.665
IND	India	2002-2019	18	2,761.8	0.513	0.396	0.583
JOR	Jordan	2002-2019	18	1,151.4	0.387	0.397	0.719
KEN	Kenya	2002-2019	18	1,596.0	0.525	0.270	0.542
KHM	Cambodia	2002-2019	18	2,878.4	0.513	0.282	0.540
LAO	Lao PDR	2002-2019	18	4,004.0	0.492	0.292	0.555
LKA	Sri Lanka	2002-2019	18	3,694.9	0.469	0.358	0.736
MAR	Morocco	2002-2019	18	1,438.9	0.392	0.376	0.612
MLI	Mali	2002-2019	18	1,380.1	0.606	0.303	0.396
MMR	Myanmar	2002-2019	18	3,626.9	0.514	0.284	0.514
MOZ	Mozambique	2002-2019	18	789.2	0.499	0.274	0.400
MWI	Malawi	2002-2019	18	1,683.8	0.556	0.288	0.454
NGA	Nigeria	2002-2019	18	1,480.9	0.500	0.319	0.493
NIC	Nicaragua	2002-2019	18	1,944.0	0.462	0.300	0.619
NPL	Nepal	2002-2019	18	2,536.2	0.514	0.324	0.543
PAK	Pakistan	2002-2019	18	2,791.7	0.530	0.359	0.508
PHL	Philippines	2002-2019	18	3,337.7	0.474	0.311	0.679

RWA	Rwanda	2002-2019	18	1,451.1	0.555	0.336	0.476
SEN	Senegal	2002-2019	18	1,160.9	0.531	0.320	0.466
THA	Thailand	2002-2019	18	3,081.8	0.434	0.424	0.743
TUN	Tunisia	2002-2019	18	1,492.5	0.372	0.419	0.716
TZA	Tanzania	2002-2019	18	1,483.9	0.520	0.303	0.488
UGA	Uganda	2002-2019	18	1,960.5	0.588	0.279	0.489
VNM	Viet Nam	2002-2019	18	5,157.8	0.477	0.368	0.661
ZMB	Zambia	2002-2019	18	2,312.7	0.493	0.318	0.521
ZWE	Zimbabwe	2002-2019	18	745.2	0.512	0.193	0.519

Figure 2 confirms that the filtered panel is balanced across the 39 retained countries and 18 feature years.



**Figure 2.** Complete country-year coverage after filtering.

The prediction target is one-year-ahead cereal yield. For each country  $c$  and feature year  $t$ , the target is  $yield(c,t+1)$ . All explanatory variables are measured at year  $t$  or earlier. The lag block includes observed yield at  $t$ , observed yield at  $t-1$ , observed yield at  $t-2$ , a three-year rolling mean, and one-year yield

growth. The macro-production block includes log cereal production, log cereal land, log GDP and a production-to-land yield check. The resilience block includes vulnerability, readiness, ND-GAIN, HDI, the resilience gap defined as vulnerability minus readiness, and an interaction between vulnerability

and low readiness. The climate block includes annual temperature, annual rainfall, temperature anomaly relative to the 1981-2000 country baseline, rainfall anomaly relative to the same baseline, and two climate-resilience interactions.

The temporal split is strict. All models are estimated on feature years 2002-2016 and evaluated on held-

out feature years 2017-2019, which correspond to target years 2018-2020. Rolling-origin validation fits the ridge panel model on all feature years before a given origin and tests on that origin. Table 3 reports the main split, and Table 4 gives descriptive statistics for the filtered panel.

**Table 3.** Time-based train/test split.

Split	Feature years	Target years	Rows	Countries
training	2002-2016	2003-2017	585	39
held-out test	2017-2019	2018-2020	117	39

**Table 4.** Descriptive statistics of the complete LMIC panel.

Variable	N	Mean	SD	Min	P10	Median	P90	Max
yield_t	702	2,366.0	1,360.9	309.7	1,091.8	1,960.7	4,193.6	7,556.2
yield_t+1	702	2,404.8	1,366.2	309.7	1,116.1	2,007.5	4,286.9	7,556.2
production	702	18.6m	44.9m	28,220.0	752,469.3	3.7m	41.1m	325.0m
land	702	6.6m	15.6m	33,497.0	433,187.8	2.3m	12.9m	101.2m
GDP	702	120.91b	306.41b	559.3m	6.19b	26.18b	299.72b	2835.61b
HDI	702	0.555	0.105	0.303	0.423	0.549	0.706	0.804
readiness	702	0.318	0.063	0.161	0.256	0.304	0.413	0.491
vulnerability	702	0.500	0.056	0.366	0.432	0.503	0.573	0.616
ND-GAIN	702	-0.182	0.098	-0.383	-0.298	-0.195	-0.035	0.064
resilience gap	702	0.182	0.098	-0.064	0.035	0.195	0.298	0.383
temperature anomaly	702	0.439	0.338	-0.625	0.087	0.395	0.849	2.537

rainfall anomaly fraction	702	0.022	0.134	-0.520	-0.129	0.017	0.184	0.663
---------------------------	-----	-------	-------	--------	--------	-------	-------	-------

Nine forecasting specifications are estimated. Persistence uses current-year yield as the next-year forecast. The country-mean baseline uses the training mean for each country. Panel fixed-effects OLS uses standardized numeric predictors and country dummy variables. Ridge adds L2 regularization to stabilize country fixed effects and correlated yield lags [19], [20]. Huber-ridge uses robust Huber loss. The nonlinear family includes a regression tree, random forest and gradient boosting model, plus residual gradient boosting that predicts the shock relative to lagged yield [21], [22]. Table 5 defines the features, and Table 6 reports the main hyperparameters.

Performance is measured with RMSE, MAE, MAPE, R2 and mean bias. RMSE and MAE are reported in kg/ha. MAPE is reported as a percentage. R2 is

computed against the held-out target mean. The consistency check compares published cereal yield with the production divided by land conversion implied by the WDI production and land indicators.

The policy-explanation step is constrained by measured diagnostics. For each selected country, the input values are vulnerability, readiness, HDI, temperature anomaly, rainfall anomaly, yield volatility, and panel medians. The generated explanation must state whether each diagnostic is above or below the panel benchmark and must select adaptation channels consistent with the measured profile. This follows the principle that model explanations should be grounded in model inputs and outputs [23], while using language-model methods only for concise policy translation [26]-[28].

**Table 5.** Feature definitions.

Feature	Definition	Source
target_yield_next	Cereal yield in kg/ha at year t+1	WDI
yield_lag1, yield_lag2, yield_lag3, yield_roll3	Observed yield at t, t-1, t-2, and their three-year mean	WDI
yield_growth1	One-year growth in cereal yield at t	WDI, constructed
log_prod, log_land	Log cereal production and log cereal land at t	WDI
log_gdp	Log GDP in current US dollars at t	WDI
vulnerability, readiness, ndgain	ND-GAIN climate exposure and adaptation-capacity scores at t	ND-GAIN
hdi	Human Development Index context variable	ND-GAIN
cckp_temp, cckp_rain	Annual average temperature and annual rainfall at t	QoG-EI / CCKP

temp_anom, rain_anom_pct	Country-specific deviations from the 1981-2000 climate baseline	constructed
resilience_gap	vulnerability - readiness	constructed
vulnerability_x_lowreadiness	vulnerability x (1 - readiness)	constructed
temp_x_lowreadiness, rainstress_x_vulnerability	climate-resilience interaction diagnostics	constructed

**Table 6.** Forecasting specifications and hyperparameters.

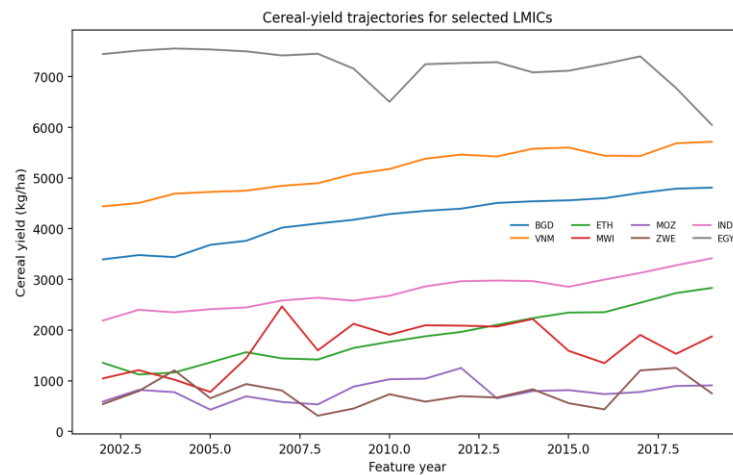
Model	Hyperparameters	Purpose
Persistence	$\hat{y}(t+1)=y(t)$	naive temporal benchmark
Country mean	country-specific mean of training target	static panel benchmark
Panel fixed-effects OLS	standardized numeric predictors plus country dummies	linear benchmark
Ridge panel	$\alpha=25$ ; standardized numeric predictors; country dummies	regularized linear panel
Huber-ridge	$\epsilon=1.35$ ; $\alpha=0.001$ ; $\max\_iter=1000$	robust linear panel
Regression tree	$\max\_depth=3$ ; $\min\_samples\_leaf=10$	single nonlinear tree
Random forest	200 trees; $\max\_depth=3$ ; $\min\_samples\_leaf=10$ ; sqrt feature subsampling	bagged nonlinear model
Gradient boosting	120 trees; $\text{learning\_rate}=0.04$ ; $\max\_depth=2$ ; $\min\_samples\_leaf=7$	additive nonlinear model
Residual gradient boosting	same boosting settings; predicts yield shock added to lagged yield	shock-forecast nonlinear model

## Results and Discussion

The empirical panel is balanced after filtering. Figure 3 illustrates yield trajectories for selected countries.

The trajectories confirm that national cereal yields are persistent but not static. Bangladesh and Viet Nam exhibit high yield levels, while several African and fragile-country cases remain much lower. This

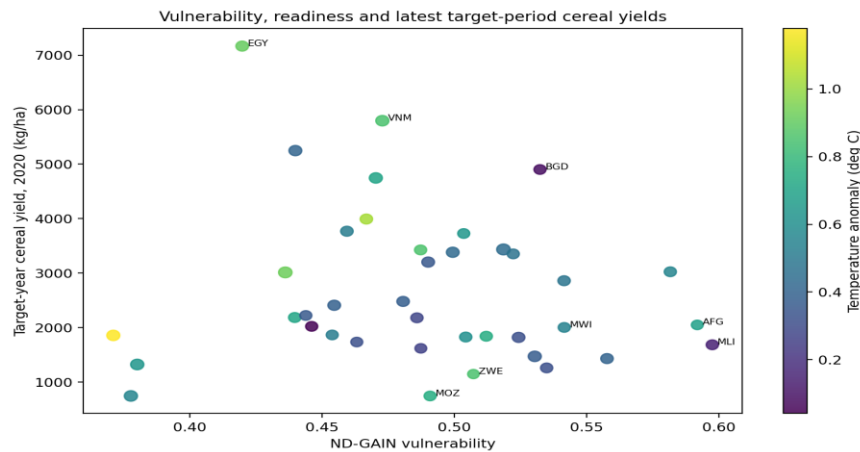
persistence motivates the inclusion of strong lag baselines and helps prevent overstatement of complex models.



**Figure 3.** Cereal-yield trajectories for selected LMICs.

Figure 4 shows that the vulnerability-yield relationship is not a one-dimensional scatter. High-vulnerability countries can have very different yield

levels because adaptation capacity, country fixed effects, crop mix and land quality mediate exposure. The figure also shows the added climate dimension by coloring countries by temperature anomaly.



**Figure 4.** ND-GAIN vulnerability, readiness and latest target-period cereal yields, with temperature anomaly shown by color.

Table 7 gives the held-out model comparison, and Figure 5 visualizes the same RMSE ranking. The ridge panel model ranks first with RMSE 318.23 kg/ha and MAE 192.93 kg/ha. Gradient boosting and Huber-ridge also perform strongly, while persistence remains a competitive benchmark. The ranking is consistent with the structure of annual national

cereal-yield data: recent yield history and country effects explain most of the short-horizon variation, while macro, resilience and climate variables provide smaller adjustments. Forecasting theory gives a similar expectation for short panels with strong temporal dependence and limited annual observations [24], [25].

**Table 7.** Held-out model comparison for target years 2018-2020.

Model	RMSE	MAE	MAPE %	R2	Bias
Ridge panel model	318.228	192.932	9.080	0.945	-26.788
Gradient boosting	323.585	199.438	9.185	0.943	-28.584
Huber-ridge panel model	325.473	198.534	9.491	0.942	-31.293
Panel fixed-effects OLS	333.216	206.991	9.323	0.939	-43.779
Random forest	345.286	238.546	11.766	0.935	-51.213
Residual gradient boosting	347.357	197.294	9.563	0.934	-21.851
Persistence (last observed yield)	355.005	202.644	9.750	0.931	-43.821
Regression tree	364.324	254.844	11.455	0.928	-47.080
Country mean baseline	505.101	389.347	15.317	0.861	-325.947

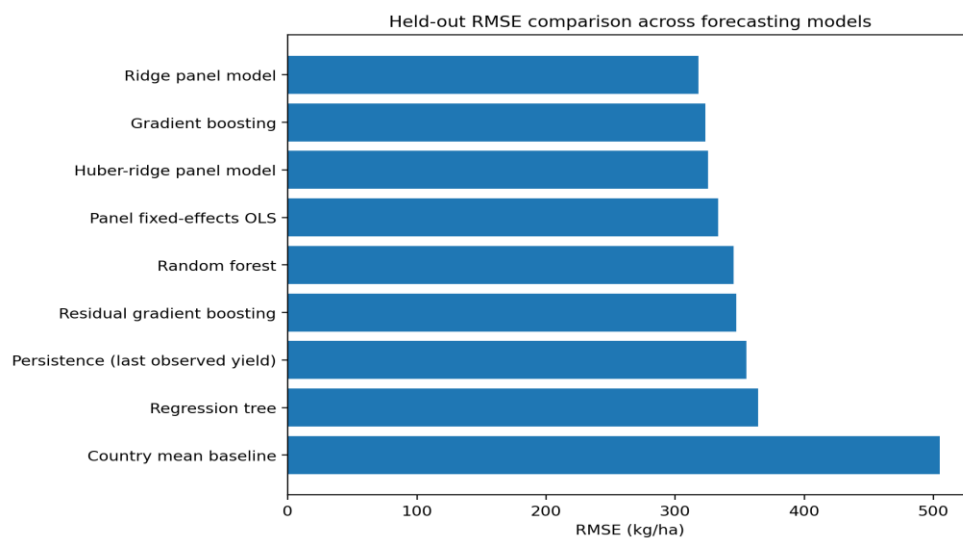
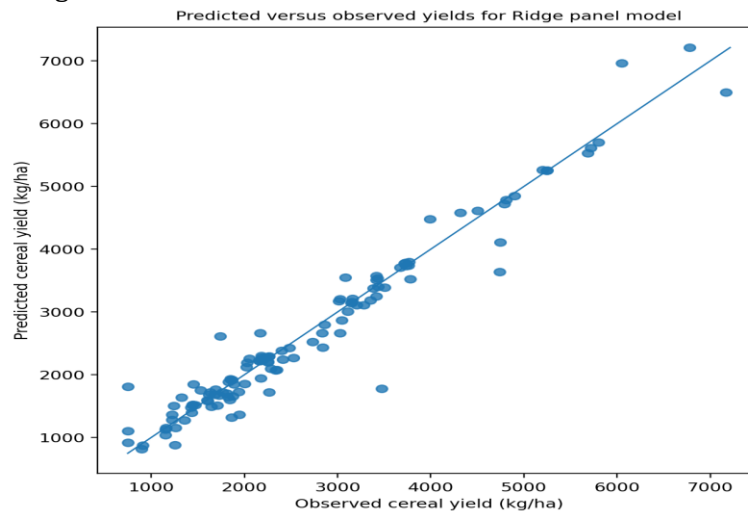
**Figure 5.** Held-out RMSE comparison across forecasting models.

Figure 6 compares predicted and observed yields for the best held-out model. Predictions cluster near the one-to-one line, with larger absolute errors

concentrated among lower-yield and more volatile country-years.



**Figure 6.** Predicted versus observed yields for the best held-out model.

Table 8 reports feature-block ablation using the same ridge panel family. The differences across the rich specifications are small. The ND-GAIN resilience block gives the lowest RMSE in the ablation table, while the full joint panel is effectively tied. The observed climate block does not dominate the one-

year forecast by itself, but it makes the climate-exposure component explicit and auditable. This is important for the paper's climate-resilience framing: ND-GAIN variables describe adaptive capacity and vulnerability, whereas annual temperature and rainfall anomalies describe measured exposure in the feature year.

**Table 8.** Feature-block ablation using the ridge panel family.

Feature block	RMSE	MAE	MAPE %	R2	Bias
+ ND-GAIN resilience	318.184	189.158	8.910	0.945	-22.552
Full joint panel	318.228	192.932	9.080	0.945	-26.788
+ observed climate	319.671	194.100	9.092	0.944	-30.436
Yield lags only	320.661	190.723	8.804	0.944	-27.165
+ WDI production/land/GDP	320.683	190.698	8.911	0.944	-29.002

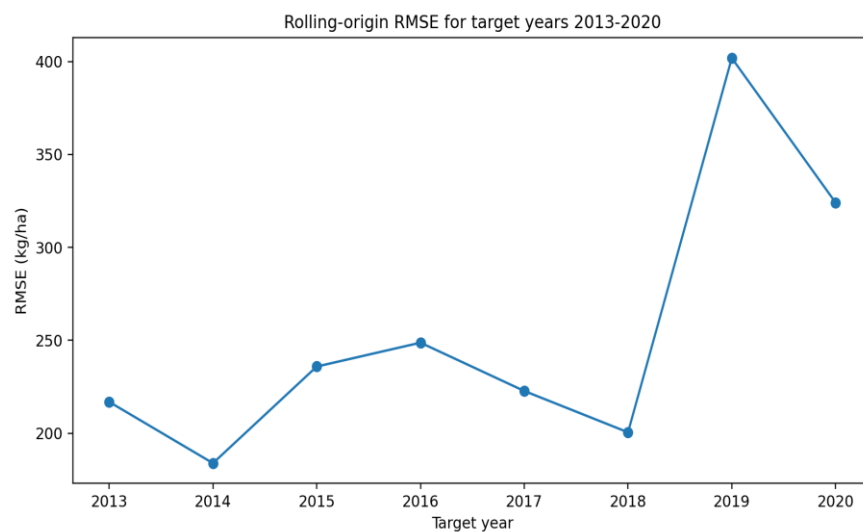
Rolling-origin validation in Table 9 and Figure 7 supports the same conclusion. The ridge panel

model remains stable across target years 2013-2020. RMSE is higher for target year 2019 than in the

surrounding origins, but all rolling R2 values remain above 0.90. The held-out 2018-2020 test is therefore not an isolated favorable split.

**Table 9.** Rolling-origin validation for the ridge panel model.

Feature year	Target year	Countries	RMSE	MAE	MAPE %	R2	Bias
2012	2013	39	216.847	138.870	9.101	0.976	44.340
2013	2014	39	183.866	116.963	6.593	0.982	9.263
2014	2015	39	235.859	165.999	9.944	0.972	45.304
2015	2016	39	248.699	144.929	10.432	0.969	98.514
2016	2017	39	222.716	153.955	7.605	0.974	22.098
2017	2018	39	200.476	151.232	6.889	0.977	-49.975
2018	2019	39	401.859	225.032	9.007	0.906	-32.275
2019	2020	39	324.086	208.309	11.687	0.947	14.264



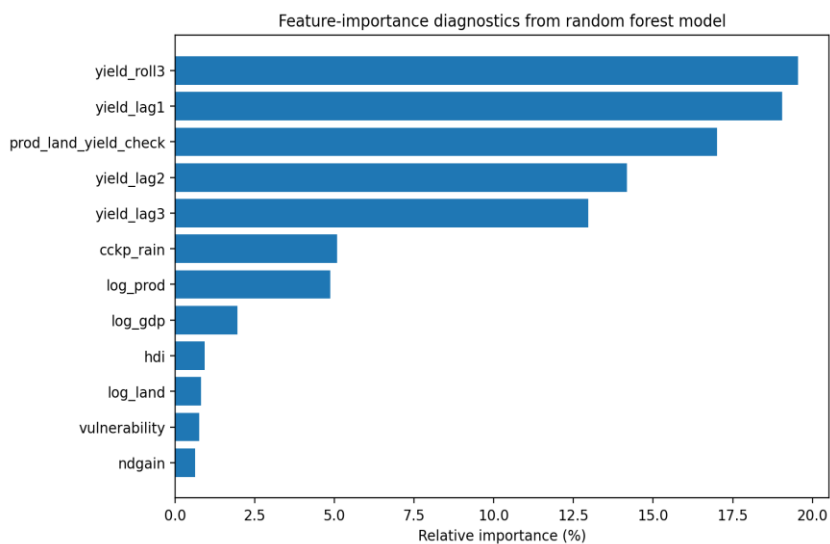
**Figure 7.** Rolling-origin RMSE for target years 2013-2020.

Table 10 and Figure 8 report model-derived feature importance from a random-forest diagnostic model. Yield lags and the production-to-land yield check account for most of the split importance. Annual rainfall also appears as a non-zero climate

diagnostic, but it is smaller than the yield-history variables. The interpretation is direct: for one-year national cereal-yield forecasting, recent yield history is the strongest predictor; for vulnerability interpretation, ND-GAIN and climate anomalies supply policy context.

**Table 10.** Feature importance from the random-forest diagnostic model.

Feature	Relative importance %
yield_roll3	19.546
yield_lag1	19.043
prod_land_yield_check	17.001
yield_lag2	14.178
yield_lag3	12.966
cckp_rain	5.085
log_prod	4.870
log_gdp	1.961
hdi	0.931
log_land	0.816
vulnerability	0.764
ndgain	0.627



**Figure 8.** Feature-importance diagnostics from the random-forest model.

Table 11 provides a consistency audit. The median absolute difference between published cereal yield

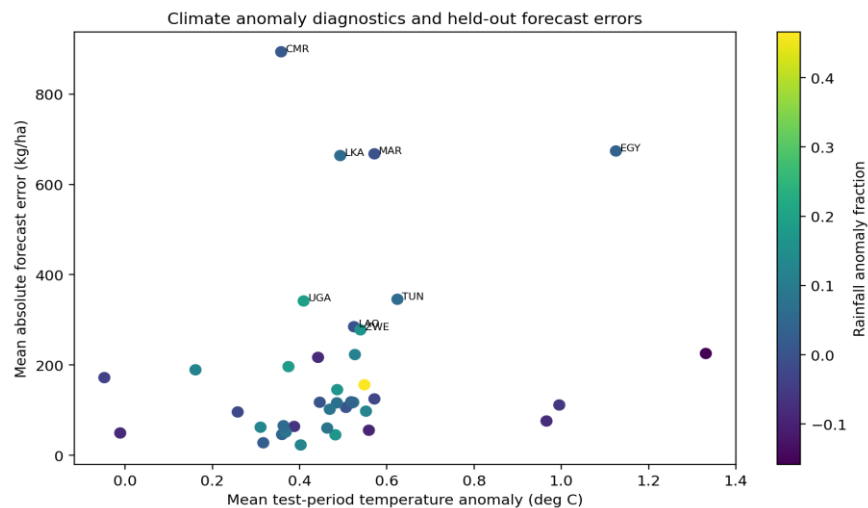
and the production-land implied yield is close to zero, and the complete-case check confirms that no

ND-GAIN or climate fields remain missing after filtering. Figure 9 relates climate anomalies to held-out forecast errors and shows why climate variables

are most useful as diagnostic context rather than as substitutes for yield history.

**Table 11.** Data consistency checks.

Check	Mean abs. diff.	Median abs. diff.	Max/check value
yield from production/land equals published cereal yield	0.025	0.026	0.050
complete ND-GAIN and climate fields after filtering			0.000



**Figure 9.** Climate anomaly diagnostics and held-out forecast errors.

The policy-explanation layer in Table 12 converts diagnostics into country-specific statements for selected vulnerable cases. Afghanistan, Mali, Uganda, Burkina Faso, Malawi and Mozambique are selected because their measured profiles combine high

vulnerability, low readiness, low HDI, positive temperature anomalies or yield volatility. The explanations do not assert unobserved weather events. They translate the measured country diagnostics into adaptation channels that a policy analyst can review.

**Table 12.** Grounded policy explanations for selected vulnerable countries.

Country	ISO3	Key diagnostics	Policy explanation
Malawi	MWI	Vuln 0.542; Ready 0.279; HDI 0.519; Temp anom. 0.547; Rain anom. 0.334; Vol. 0.210	Malawi shows elevated adaptation risk because ND-GAIN vulnerability is above the panel median (0.542) and readiness is

			below the median (0.279). HDI is below the median (0.519), while the 2019 climate diagnostics are warmer than its 1981-2000 country baseline and wetter than or close to its 1981-2000 rainfall baseline. The measured held-out yield volatility is 0.210. Policy attention should focus on water-management infrastructure, resilient seed systems, grain storage, and extension support in cereal-producing districts.
Afghanistan	AFG	Vuln 0.592; Ready 0.249; HDI 0.488; Temp anom. 0.681; Rain anom. 0.149; Vol. 0.053	Afghanistan shows elevated adaptation risk because ND-GAIN vulnerability is above the panel median (0.592) and readiness is below the median (0.249). HDI is below the median (0.488), while the 2019 climate diagnostics are warmer than its 1981-2000 country baseline and wetter than or close to its 1981-2000 rainfall baseline. The measured held-out yield volatility is 0.053. Policy attention should focus on water-management infrastructure, resilient seed systems, grain storage, and extension support in cereal-producing districts.
Mali	MLI	Vuln 0.598; Ready 0.297; HDI 0.433; Temp anom. 0.144; Rain anom. 0.173; Vol. 0.027	Mali shows elevated adaptation risk because ND-GAIN vulnerability is above the panel median (0.598) and readiness is

below the median (0.297). HDI is below the median (0.433), while the 2019 climate diagnostics are warmer than its 1981-2000 country baseline and wetter than or close to its 1981-2000 rainfall baseline. The measured held-out yield volatility is 0.027. Policy attention should focus on water-management infrastructure, resilient seed systems, grain storage, and extension support in cereal-producing districts.

Burkina Faso

BFA

Vuln 0.535; Ready 0.291;  
HDI 0.452; Temp anom.  
0.321; Rain anom. 0.213;  
Vol. 0.065

Burkina Faso shows elevated adaptation risk because ND-GAIN vulnerability is above the panel median (0.535) and readiness is below the median (0.291). HDI is below the median (0.452), while the 2019 climate diagnostics are warmer than its 1981-2000 country baseline and wetter than or close to its 1981-2000 rainfall baseline. The measured held-out yield volatility is 0.065. Policy attention should focus on water-management infrastructure, resilient seed systems, grain storage, and extension support in cereal-producing districts.

Uganda

UGA

Vuln 0.582; Ready 0.287;  
HDI 0.525; Temp anom.  
0.559; Rain anom. 0.323;  
Vol. 0.010

Uganda shows elevated adaptation risk because ND-GAIN vulnerability is above the panel median (0.582) and readiness is

---

below the median (0.287). HDI is below the median (0.525), while the 2019 climate diagnostics are warmer than its 1981-2000 country baseline and wetter than or close to its 1981-2000 rainfall baseline. The measured held-out yield volatility is 0.010. Policy attention should focus on water-management infrastructure, resilient seed systems, grain storage, and extension support in cereal-producing districts.

---

Mozambique

MOZ

Vuln 0.491; Ready 0.270;  
HDI 0.456; Temp anom.  
0.744; Rain anom. 0.160;  
Vol. 0.065

Mozambique shows elevated adaptation risk because ND-GAIN vulnerability is at or below the panel median (0.491) and readiness is below the median (0.270). HDI is below the median (0.456), while the 2019 climate diagnostics are warmer than its 1981-2000 country baseline and wetter than or close to its 1981-2000 rainfall baseline. The measured held-out yield volatility is 0.065. Policy attention should focus on water-management infrastructure, resilient seed systems, grain storage, and extension support in cereal-producing districts.

---

## Limitations

The first limitation is crop aggregation. The empirical target is the harmonized cereal-yield aggregate, which supports consistent country coverage and direct consistency with cereal production and land indicators. The aggregate is appropriate for national food-security forecasting, but it does not identify crop-specific mechanisms for wheat, rice and maize separately. A crop-specific extension should add separate FAOSTAT item codes and repeat the panel construction where complete country-year coverage exists.

The second limitation is spatial scale. Country-year annual data are useful for policy screening, but they hide subnational climate gradients, irrigation heterogeneity, conflict disruption and market-access differences. Subnational crop maps, gridded weather variables and remote-sensing vegetation indices would improve district-level targeting. They were not added here because the objective is a compact national panel built from public country-year releases.

The third limitation is causal interpretation. The models forecast next-year yield; they do not estimate causal effects of readiness, vulnerability, GDP, temperature or rainfall on yields. ND-GAIN variables summarize exposure and adaptive capacity, and they change slowly. Their modest direct contribution to one-year accuracy does not imply that adaptation policy is unimportant. It means that recent yield history dominates short-horizon national forecasts, while resilience variables are better suited to prioritization and explanation.

The fourth limitation is the policy-explanation layer. The text is grounded in measured diagnostics and constrained language, but it remains a communication layer. It should be reviewed by country specialists before operational use. A production system should add document retrieval, human review and jurisdiction-specific policy constraints before recommendations are used in program design.

## Conclusion

This paper conducted an empirical evaluation of climate-resilience and cereal-yield joint forecasting for 39 LMICs. The final panel contains 702 complete country-year rows and combines WDI cereal yield, WDI cereal production, land and GDP indicators, ND-GAIN vulnerability, readiness, ND-GAIN and HDI variables, and QoG-EI annual temperature and rainfall variables. The held-out test covers target years 2018-2020, and the best model is a ridge panel specification with RMSE 318.23 kg/ha, MAE 192.93 kg/ha and R<sup>2</sup> 0.945.

The core empirical finding is that recent yield history dominates short-horizon country-level cereal-yield forecasting, while resilience and climate-exposure indicators add their clearest value in explanation, diagnostics and prioritization. Forecasting and vulnerability explanation should be linked, but they should not be confused. The forecast model should remain accountable to out-of-sample metrics, and the policy-explanation layer should remain grounded in measured vulnerability, readiness, HDI, climate anomalies and yield-volatility diagnostics.

## References

- [1] Intergovernmental Panel on Climate Change, *Climate Change 2022: Impacts, Adaptation and Vulnerability*. Cambridge, U.K.: Cambridge Univ. Press, 2022.
- [2] Food and Agriculture Organization of the United Nations, *World Food and Agriculture: Statistical Yearbook 2022*. Rome, Italy: FAO, 2022.
- [3] World Bank, *World Development Indicators*, December 16, 2024 Archive. Washington, DC, USA: World Bank, 2024.
- [4] Notre Dame Global Adaptation Initiative, *ND-GAIN Country Index Technical Report*. Notre Dame, IN, USA: Univ. of Notre Dame, 2015.
- [5] Quality of Government Institute, *QoG Environmental Indicators Dataset*. Gothenburg, Sweden: Univ. of Gothenburg, 2024.
- [6] I. C. Harris, T. J. Osborn, P. Jones, and D. Lister, "Version 4 of the CRU TS monthly high-resolution gridded multivariate climate dataset," *Scientific Data*, vol. 7, 2020.
- [7] D. B. Lobell and C. B. Field, "Global scale climate-crop yield relationships and the impacts of

- recent warming," *Environmental Research Letters*, vol. 2, no. 1, 2007.
- [8] W. Schlenker and M. J. Roberts, "Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change," *Proceedings of the National Academy of Sciences*, vol. 106, no. 37, pp. 15594-15598, 2009.
- [9] D. B. Lobell, W. Schlenker, and J. Costa-Roberts, "Climate trends and global crop production since 1980," *Science*, vol. 333, no. 6042, pp. 616-620, 2011.
- [10] A. J. Challinor et al., "A meta-analysis of crop yield under climate change and adaptation," *Nature Climate Change*, vol. 4, pp. 287-291, 2014.
- [11] T. Wheeler and J. von Braun, "Climate change impacts on global food security," *Science*, vol. 341, no. 6145, pp. 508-513, 2013.
- [12] D. K. Ray, N. D. Mueller, P. C. West, and J. A. Foley, "Yield trends are insufficient to double global crop production by 2050," *PLoS ONE*, vol. 8, no. 6, 2013.
- [13] D. K. Ray, J. S. Gerber, G. K. MacDonald, and P. C. West, "Climate variation explains a third of global crop yield variability," *Nature Communications*, vol. 6, 2015.
- [14] C. Lesk, P. Rowhani, and N. Ramankutty, "Influence of extreme weather disasters on global crop production," *Nature*, vol. 529, pp. 84-87, 2016.
- [15] C. Rosenzweig et al., "Assessing agricultural risks of climate change in the 21st century in a global gridded crop model intercomparison," *Proceedings of the National Academy of Sciences*, vol. 111, no. 9, pp. 3268-3273, 2014.
- [16] S. J. Vermeulen, B. M. Campbell, and J. S. I. Ingram, "Climate change and food systems," *Annual Review of Environment and Resources*, vol. 37, pp. 195-222, 2012.
- [17] J. Knox, T. Hess, A. Daccache, and T. Wheeler, "Climate change impacts on crop productivity in Africa and South Asia," *Environmental Research Letters*, vol. 7, no. 3, 2012.
- [18] M. Auffhammer, S. M. Hsiang, W. Schlenker, and A. Sobel, "Using weather data and climate model output in economic analyses of climate change," *Review of Environmental Economics and Policy*, vol. 7, no. 2, pp. 181-198, 2013.
- [19] B. H. Baltagi, *Econometric Analysis of Panel Data*, 3rd ed. Chichester, U.K.: Wiley, 2005.
- [20] A. E. Hoerl and R. W. Kennard, "Ridge regression: Biased estimation for nonorthogonal problems," *Technometrics*, vol. 12, no. 1, pp. 55-67, 1970.
- [21] Jiaying Jin, Tina Huang, and Sam Lu, "A Model-Risk-Friendly Probability of Default Workflow: Calibration, Distribution-Free Uncertainty Quantification, and SHAP Explanations on the UCI Credit Card Default Dataset", *JACS*, vol. 4, no. 6, pp. 74-85, Jun. 2024, doi: 10.69987/JACS.2024.40606.
- [22] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *Annals of Statistics*, vol. 29, no. 5, pp. 1189-1232, 2001.
- [23] Chenyu Li, Wenhao Su, and Eric Zhang, "Lightweight Hallucination Firewall for Enterprise LLM Applications: Evidence Consistency, Self-Checking, and Small-Model Detection on TruthfulQA", *JACS*, vol. 3, no. 1, pp. 49-65, Jan. 2023, doi: 10.69987/JACS.2023.30104.
- [24] R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*, 3rd ed. Melbourne, Australia: OTexts, 2021.
- [25] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning*, 2nd ed. New York, NY, USA: Springer, 2021.
- [26] A. Vaswani et al., "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017.
- [27] Xiaohan Chang, Tong Ye, and Sophia Luo, "LLM-as-Reranker for Personalized Recommendation: Popularity Bias Mitigation and Faithful Natural-Language Explanations on MovieLens 100K", *JACS*, vol. 3, no. 8, pp. 61-78, Aug. 2023, doi: 10.69987/JACS.2023.30806.
- [28] L. Ouyang et al., "Training language models to follow instructions with human feedback," in *Advances in Neural Information Processing Systems*, 2022.