Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Forecasting agricultural value added under drought risk in Morocco: a hybrid Fourier ARDL–neural network approach

Lamiaa Aattar

lamiaa.aattar@um5s.net.ma

Mohammed V University

Saad Elouardirhi

Mohammed V University of Rabat

Research Article

Keywords: Forecasting, Fourier ARDL, Neural network, Hybrid model, Drought, Agricultural value added, Morocco

Posted Date: November 18th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11515411/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Forecasting agricultural value added under drought risk in Morocco: a hybrid Fourier ARDL–neural network approach

Lamiaa Aattar[1*] and Saad Elouardirhi[1]

[1*]Faculty of Legal, Economic and Social Sciences Souissi, Mohammed V University of Rabat, Rabat, Morocco.

*Corresponding author(s). E-mail(s): lamiaa.aatar@gmail.com;

Abstract

Droughts have become longer and more frequent in Morocco, and their effect on agricultural output is now a major source of uncertainty in macroeconomic forecasts. This paper asks how well agricultural value added can be forecast once crop-year rainfall is known, and what droughts of different intensity and duration would cost. We build a consistent quarterly series of agricultural value added for 1998–2025, benchmarked to the latest national accounts, and link it to rainfall and temperature from the ERA5 reanalysis. We propose a hybrid model in which a Fourier ARDL captures the linear relation and gradual structural change, and a small neural network, selected on a separate validation period, models the nonlinear structure left in its residuals. A 10 percent fall in crop-year rainfall lowers agricultural value added by about 4.4 percent in the long run. Over 2020–2025, the Fourier terms improve four-quarter-ahead forecasts by about 30 percent relative to a standard ARDL, and the hybrid model is the most accurate of six models, with a mean absolute percentage error of 5.4 percent against 11.1 percent for ARIMA, although the network's own contribution is modest. Scenario simulations show that a two-year drought costs more than a single severe one, about 55 against 41 billion dirhams.

Keywords: Forecasting, Fourier ARDL, Neural network, Hybrid model, Drought, Agricultural value added, Morocco

JEL Classification: C22 , C45 , C53 , Q10 , Q54

1 Introduction

Every spring, Moroccan forecasters face the same question. By the end of April the rainy season is over, the cereal crop is largely made, and the amount of rain that fell since October is known with some precision. What is not yet known is what this rain will mean for agricultural output, and through it for growth, employment, imports of grain and the budget. The answer matters a great deal. Agriculture accounts for a little over 10 percent of GDP and close to 28 percent of employment (World Bank, 2026), and a single poor campaign is enough to pull national growth down, as in 2022 when GDP grew by only 1.5 percent. The reverse is also true. After the abundant rains of the 2025/26 season, the Ministry of Agriculture expects agricultural GDP to rise by around 15 percent.

Forecasting this sector has become harder as droughts have become more frequent and longer. Between 2022 and 2024, Morocco went through three consecutive campaigns with rainfall well below normal, the longest dry sequence of the last three decades. These are exactly the years in which reliable forecasts are most needed, and in which models built on calmer periods are most likely to fail.

Existing work on Morocco offers useful pieces of an answer, but none that fits this question directly. A first group of studies forecasts crop yields. Balaghi et al (2008) used rainfall, temperature and satellite vegetation indices to predict wheat yields at the provincial and national level, and Bouras et al (2021) showed that machine-learning methods such as XGBoost, combined with drought indices from remote sensing, can predict cereal yields four months before harvest. These studies are agronomic. They work with yields in tonnes per hectare, at the province level, and do not say how a drought translates into value added, the variable that enters the national accounts and macroeconomic projections.

A second group forecasts macroeconomic aggregates. Daoui (2023) applied factor-augmented error-correction models to more than a hundred Moroccan series, and a recent IMF working paper (International Monetary Fund, 2026) used satellite indices, Google Trends and flight data to nowcast several indicators, including agricultural value added. Nowcasting, however, estimates the current quarter before official figures are released. It is not designed to forecast a year ahead, and it does not let the user ask what would happen under a different rainfall path.

A third group, closest to ours in method, combines statistical models with neural networks. Since Zhang (2003), hybrid models in which a linear model captures the main structure and a neural network is fitted to its residuals have been applied to many series, and they often forecast better than either component alone (Khashei and Bijari, 2011). In Morocco, Soumbara et al (2025) recently used a SARIMAX–multilayer perceptron model to forecast electricity demand under heat waves. We are not aware of any study that applies this approach to Moroccan agriculture.

There is also a more basic difficulty. Most hybrid models use ARIMA or SARIMAX as their linear part. These models are built for forecasting, not for interpretation: once the series has been differenced, they say little about the long-run effect of climate on output. They also assume that the relation they estimate is stable. Moroccan agriculture has changed a great deal over the past twenty-five years, with the Plan Maroc Vert, the spread of drip irrigation and repeated droughts, and these changes

have been gradual rather than abrupt. Fourier approximations, which capture such smooth shifts with a pair of trigonometric terms, are now common in cointegration analysis (Banerjee et al, 2017; Yilanci et al, 2020), but they have rarely been used for forecasting. Whether they improve forecasts is an open question.

This paper brings these strands together. We build a quarterly series of Moroccan agricultural value added for 1998–2025, link it to crop-year rainfall and spring temperature from the ERA5 reanalysis, and propose a hybrid model in which a Fourier ARDL forms the linear part and a neural network models what is left in its residuals. We then use the model to simulate the cost of droughts of different intensity and duration over 2026–2028.

The paper makes three contributions. The first is methodological. To our knowledge, it is the first to use a Fourier ARDL as the linear component of a hybrid forecasting model. This keeps the economic content of the ARDL, with short-run and long-run rainfall elasticities and a speed of adjustment, while letting the network pick up nonlinear patterns that the linear model misses. We do not add the network by default: we test the ARDL residuals for nonlinearity first, choose the network's architecture on a validation period separate from the test period, and compare the hybrid with five benchmarks, including a hybrid ARIMAX–neural network model built the same way.

The second contribution concerns the data. The Haut-Commissariat au Plan (HCP) has changed its base year twice since 1998, and its quarterly and annual figures do not always agree. We build a consistent quarterly series by benchmarking the quarterly figures to the latest annual accounts, and we define rainfall over the crop year, from October to April, rather than over the calendar year. Both choices matter for the results, and we document them so that others can reuse the series.

The third contribution is a set of drought scenarios with measures of uncertainty. We translate rainfall deficits into losses of agricultural value added, in percent and in dirhams, and we compare a single severe drought with a prolonged one of the kind Morocco experienced in 2022–2024.

The results can be summarised briefly. Rainfall has a large and stable effect: in the long run, a 10 percent fall in crop-year rainfall lowers agricultural value added by about 4.4 percent. The Fourier terms are significant and improve forecasts substantially, by around 30 percent at a four-quarter horizon compared with a standard ARDL. The neural network improves the forecasts further at this horizon, and the hybrid model is the most accurate of the six models we compare, with an error half that of ARIMA. The gain from the network is modest, though, and we discuss why. The scenarios show that the duration of a drought matters as much as its intensity: two dry years at 40 percent below normal cost more, in total, than a single year at 50 percent below normal. The effect of temperature, by contrast, cannot be identified reliably with national data, a limitation we discuss openly.

The rest of the paper is organised as follows. Section 2 reviews the literature. Section 3 describes the data and the construction of the variables. Section 4 presents the methodology. Section 5 reports the results, and Section 6 concludes with policy implications.

2 Literature review

2.1 Climate and Moroccan agriculture

That rainfall drives Moroccan agricultural output is not in doubt. Most cereals are grown without irrigation, and the country's climate combines a short rainy season with large swings from one year to the next. Schilling et al (2012) describe Morocco as one of the North African countries most exposed to climate change, with projections pointing to less rain and higher temperatures. Driouech et al (2021) show that these trends are already visible in observed data, with a clear warming across the country and drier conditions in several regions.

The economic side of this exposure has mostly been studied with cointegration methods on annual data. Aattar and Elouardirhi (2022) examined the contribution of agriculture to Moroccan growth with a cointegration analysis, and a growing number of studies use ARDL models to relate agricultural output to rainfall and energy use. This line of work explains the past well. They were not built to forecast, and annual data, with fewer than fifty observations, leave little room for the more flexible models that forecasting often requires.

The forecasting literature on Moroccan agriculture comes mainly from agronomy and remote sensing. Balaghi et al (2008) predicted wheat yields with regressions on rainfall, temperature and vegetation indices, and showed that useful national forecasts can be made from March onward. Bouras et al (2021) extended this work with machine-learning methods and satellite drought indices for the fifteen main cereal provinces. Their best model, XGBoost, explained close to 90 percent of yield variance four months before harvest. Two features of their results are worth noting for our purposes. Weather data were the most useful inputs for early forecasts, and satellite indices only became decisive close to harvest. This supports the idea that a model based on crop-year rainfall can provide early and reasonably accurate signals.

2.2 Gradual structural change and the Fourier approach

Economic relations that hold over twenty-five years rarely stay fixed. The usual answer is to add dummy variables for known break dates, but this requires knowing when the breaks happened and assumes they were sudden. For a sector shaped by policies that took years to spread, such as subsidies for drip irrigation under the Plan Maroc Vert, neither assumption is convincing.

Becker et al (2006) and Enders and Lee (2012) proposed instead to approximate breaks of unknown number, timing and form with a small number of sine and cosine terms. A single frequency is often enough to capture smooth shifts in the mean, and the cost in degrees of freedom is low. The idea was carried over to cointegration by Banerjee et al (2017), who developed a Fourier version of the ADL test, and combined with the ARDL bounds test, together with bootstrap critical values, by McNown et al (2018) and Yilanci et al (2020). Since then, Fourier ARDL models have been widely used in energy and environmental economics. Their use for forecasting is much less common, and we have not found an application in which a Fourier ARDL serves as the linear component of a hybrid model.

2.3 Hybrid statistical and neural network models

Zhang (2003) started from a simple observation: real series often contain both linear and nonlinear patterns, and no single model handles both well. ARIMA models miss nonlinear dynamics, while neural networks can be unstable when the structure is mostly linear. His solution was to fit an ARIMA model first and then a neural network to its residuals, adding the two forecasts. Variants of this scheme have been applied to stock prices (Pai and Lin, 2005), water quality (Faruk, 2010), fuelwood prices (Koutroumanidis et al, 2009) and wind speed (Cadenas and Rivera, 2010), and Khashei and Bijari (2011) proposed a more flexible combination of the two components. In Morocco, Soumbara et al (2025) built a SARIMAX–multilayer perceptron model for electricity demand with temperature as an input, and showed that it captured demand surges during simulated heat waves better than SARIMA or SARIMAX alone.

The evidence is not uniformly favourable, however. Taskaya-Temizel and Casey (2005) found that hybrids do not always beat their components and can do worse when the residuals contain little nonlinear structure. More broadly, Makridakis et al (2018) showed on a large set of series that machine-learning methods often forecast worse than simple statistical models when the series are short. Two lessons follow. The case for a hybrid should be tested rather than assumed, by checking the residuals for nonlinearity before fitting a network. And the network should be kept small, with its architecture chosen out of sample, if it is not to overfit a sample of a hundred observations.

2.4 Where this paper fits

Putting these strands together, three gaps stand out. Studies of climate and Moroccan agriculture based on cointegration explain the past but do not forecast. Forecasting studies either work with crop yields at the province level, which do not map directly into national accounts, or with macroeconomic nowcasting, which does not allow scenarios. And hybrid models, including the one recent Moroccan application, use ARIMA-type linear components that discard the long-run economic content and assume stable parameters.

Our paper addresses the three at once. We forecast agricultural value added itself, the variable used in macroeconomic and budget projections. We use a Fourier ARDL as the linear component, so that the model keeps interpretable rainfall elasticities while allowing for gradual change. And we follow the cautions of the literature by testing for nonlinearity before adding the network, selecting it on a validation period, and comparing it against a hybrid ARIMAX built in the same way.

3 Data

3.1 Agricultural value added

The dependent variable is the value added of the agriculture branch, measured in volume (previous-year prices, chain-linked) and seasonally adjusted, as published in the quarterly national accounts of the HCP. We use the log of this series. The sample runs from 1998Q1 to 2025Q4, which gives 112 quarterly observations.

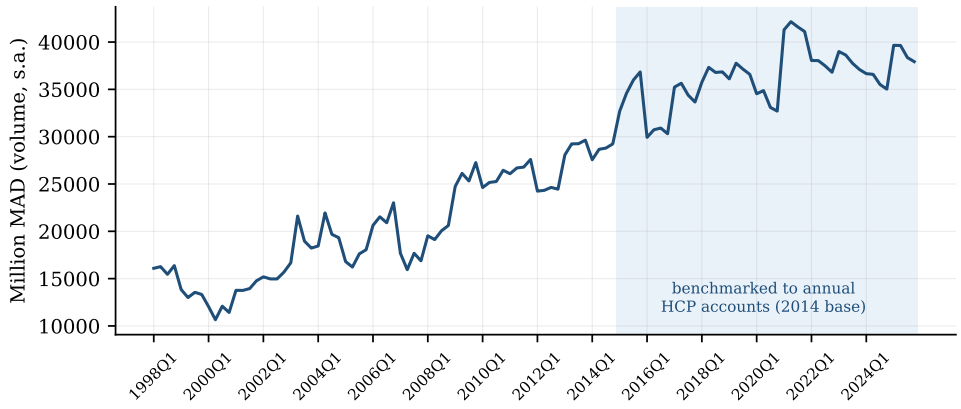


Fig. 1 Quarterly agricultural value added in Morocco, 1998Q1–2025Q4
Source: HCP quarterly and annual national accounts; author's calculations. The shaded area marks quarters benchmarked to annual accounts in the 2014 base with the proportional Denton method.

Building a long and consistent quarterly series for Morocco is less simple than it sounds. HCP has changed its base year twice over the period (1998, 2007 and 2014), and the quarterly figures for 1998 to 2014 are only available in the 1998 base. For the years after 2014 we have the year-on-year quarterly growth rates in the 2014 base, but these do not always add up to the annual growth rates that HCP publishes later in its revised accounts. The gap is large in some years. For 2015, the quarterly figures imply growth of about 14 percent, while the final annual accounts report 22.6 percent.

We dealt with this in two steps. First, we kept the 1998–2014 quarterly levels unchanged and extended the series forward with the quarterly growth rates. Second, we benchmarked the 2015–2025 quarters to the annual growth rates of the latest HCP vintage, using the proportional Denton method (Denton, 1971). Denton keeps the quarterly profile of the indicator as close as possible to the original while forcing each year to match its annual total. After benchmarking, the series reproduces every annual growth rate published by HCP between 2015 and 2025, including +1.4 percent in 2023, −5.7 percent in 2024 and +8.2 percent in 2025. Figure 1 shows the resulting series.

One feature of the data matters for the rest of the paper. Within a given year, the four quarterly growth rates are almost identical. HCP spreads the outcome of the crop year across the quarters of the calendar year, so the quarterly series carries an annual agricultural signal rather than a true intra-year cycle. This is why climate enters the model at the level of the crop year, and not quarter by quarter.

3.2 Climate variables

Climate data come from the ERA5 reanalysis of the European Centre for Medium-Range Weather Forecasts (Hersbach et al, 2020), at a 0.25-degree resolution, averaged over the whole national territory and downloaded at monthly frequency from the World Bank Climate Change Knowledge Portal. We preferred ERA5 to CRU TS because it is updated more quickly and covers the full 2022–2025 drought period.

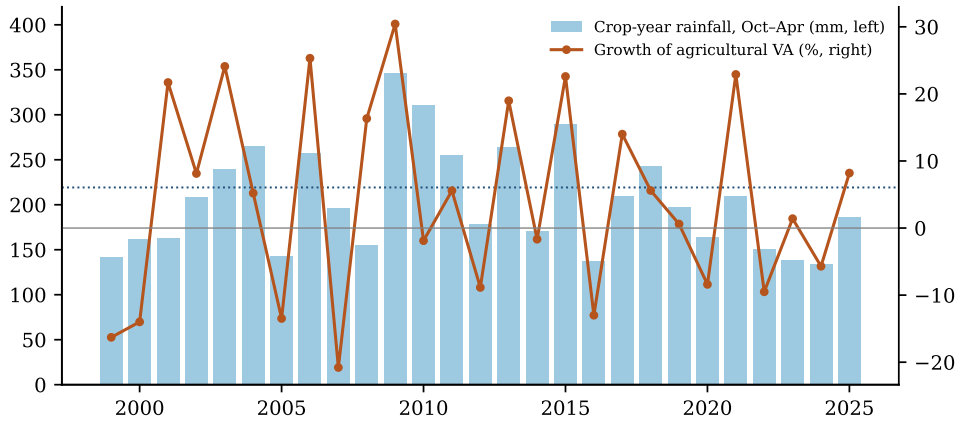


Fig. 2 Crop-year rainfall and growth of agricultural value added, 1999–2025
Source: ERA5 via the World Bank Climate Change Knowledge Portal; HCP; author's calculations. The dotted line is the 1991–2020 rainfall norm (219 mm).

Moroccan rainfed agriculture depends on the rain that falls between October and April. We therefore define *crop-year rainfall* as total precipitation from October of year $t-1$ to April of year t, and we assign it to the four quarters of year t, which is consistent with the way HCP distributes the agricultural year. A wider September–May window, often used in campaign monitoring, gives similar elasticities but a poorer fit on the quarterly series; we report it as a robustness check in Section 5.6. Over 1991–2020, crop-year rainfall averaged 219 mm. The three campaigns of 2022, 2023 and 2024 were between 31 and 39 percent below this norm, the longest dry sequence in our sample.

For temperature we use the mean daily maximum temperature from March to May, the months when cereals flower and fill their grain and are most exposed to heat stress. Two alternatives, the mean temperature and the mean maximum temperature over the whole crop year, are used to test the robustness of the results.

Because the climate variables change only once a year, their time-series properties are tested on the 28 annual observations rather than on the quarterly data, where the repeated values would inflate the sample artificially.

The simple correlation between crop-year rainfall and the annual growth of agricultural value added is 0.60 over 1999–2025 (Figure 2). The dry years of 2007, 2016, 2020 and 2022 all coincide with sharp falls in output. There are exceptions, though. In 2023, rainfall was 37 percent below normal, yet value added rose slightly, partly as a rebound from the collapse of 2022. Episodes like this suggest that the link between rain and output is not purely linear, which motivates the hybrid approach.

4 Methodology

4.1 Order of integration

We first check that no series is integrated of order two, since the ARDL bounds test is not valid for I(2) variables. We apply the ADF, Phillips–Perron and KPSS tests, with a constant and with a constant and trend. Because standard tests can confuse an unmodelled structural change with a unit root, we also apply the Fourier ADF test of Enders and Lee (2012). This test approximates smooth breaks of unknown number and form with a single pair of sine and cosine terms:

$$\Delta y_t = \rho\, y_{t-1} + c_0 + c_1 t + \gamma_1 \sin\!\left(\frac{2\pi k t}{T}\right) + \gamma_2 \cos\!\left(\frac{2\pi k t}{T}\right) + \sum_{j=1}^{p} \phi_j \Delta y_{t-j} + \varepsilon_t, \quad (1)$$

where k is the frequency, selected among 1 to 5 by minimising the sum of squared residuals, and T is the sample size.

4.2 The Fourier ARDL model

The linear part of the hybrid model is an autoregressive distributed lag model augmented with Fourier terms, following Yilanci et al (2020). In error-correction form, it reads:

$$\begin{aligned}\Delta \ln VA_t = \alpha_0 + \alpha_1 t + \theta \ln VA_{t-1} + \beta_1 \ln R_{t-1} + \beta_2 TX_{t-1} + \sum_{i=0}^{q-1} \delta_i \Delta \ln R_{t-i} + \sum_{i=0}^{r-1} \lambda_i \Delta TX_{t-i} \\ + \gamma_1 \sin\!\left(\frac{2\pi k t}{T}\right) + \gamma_2 \cos\!\left(\frac{2\pi k t}{T}\right) + u_t, \quad (2)\end{aligned}$$

where VA is agricultural value added, R is crop-year rainfall and TX is spring maximum temperature. The long-run elasticities are given by $-\beta_1/\theta$ and $-\beta_2/\theta$, and θ measures the speed at which output returns to its long-run path after a shock.

The ARDL framework suits our data for three reasons. It accepts a mix of I(0) and I(1) regressors, which is what the unit root tests show. It separates short-run from long-run effects of rainfall, something an ARIMAX model cannot do once the dependent variable has been differenced. And it performs reasonably well in samples of the size we have. The Fourier terms allow the intercept to shift gradually, which is a plausible description of a sector that went through the Plan Maroc Vert, the expansion of drip irrigation and repeated droughts over the period.

Lag orders are chosen by the Akaike criterion with a maximum of four lags. The frequency k is chosen by comparing the models for $k = 0$ to 3, and the Fourier terms are kept only if they are jointly significant in the ARDL. Cointegration is tested with the bounds test of Pesaran et al (2001), in its case with an unrestricted constant and trend. The model is then checked for serial correlation (Breusch–Godfrey), heteroskedasticity (Breusch–Pagan, ARCH), functional form (Ramsey RESET) and parameter stability (CUSUM).

4.3 The hybrid Fourier ARDL–ANN model

Following Zhang (2003), we treat the series as the sum of a linear and a nonlinear component, $\ln VA_t = L_t + N_t$. The Fourier ARDL gives an estimate of the linear part, $\hat{L}_t$. Its residuals, $e_t = \ln VA_t - \hat{L}_t$, should contain no linear structure if the model is well specified, but they may still carry nonlinear patterns. We test for this with the BDS test of Broock et al (1996), applied to the ARDL residuals. Only if the test rejects independence do we fit a neural network to the residuals.

The nonlinear part is estimated with a multilayer perceptron with one hidden layer and a hyperbolic tangent activation, $\hat{N}_t = f(e_{t-1}, \ldots, e_{t-p})$, and the hybrid forecast is $\widehat{\ln VA}_t = \hat{L}_t + \hat{N}_t$.

Small samples make neural networks easy to overfit, and we took two precautions. First, the architecture was not chosen on the test period. We set aside 2016Q1–2019Q4 as a validation period and compared eighteen configurations (one, two or four lags of the residuals, with or without the climate variables as extra inputs, and three, five or ten hidden neurons). The configuration with the lowest validation error was retained and then frozen. Second, because the weights of a single network depend on their random starting values, each forecast is the median of fifteen networks trained from different seeds.

4.4 Forecast evaluation

The models are estimated on 1998Q1–2019Q4 and evaluated on 2020Q1–2025Q4. The test period is demanding: it contains the drought of 2020, the strong recovery of 2021, three consecutive dry years and the partial rebound of 2025. We use an expanding window. At each step the model is re-estimated with all data available up to that point, a forecast is made, and the window moves forward one quarter. Forecasts are produced one quarter ahead and four quarters ahead. At the four-quarter horizon the forecasts are recursive, with predicted values of agricultural output fed back into the lagged terms.

All forecasts are conditional on the observed climate of the crop year. This is not an unrealistic assumption in the Moroccan setting. The rainfall of a campaign is known at the end of April, well before the harvest is completed and long before HCP publishes the national accounts for the year. The question our forecasts answer is therefore the one that the Ministry of Economy and Finance and Bank Al-Maghrib face every spring: given the rain we have had, what will agricultural output be?

The hybrid model is compared with five benchmarks: ARIMA, ARIMAX with the same climate variables, a hybrid ARIMAX–ANN built in the same way as our model, a standard ARDL without Fourier terms, and the Fourier ARDL alone. ARIMA and ARIMAX orders are chosen once by AIC on the estimation sample. Accuracy is measured by the root mean squared error, the mean absolute error and the mean absolute percentage error, computed on the level of value added. Differences in accuracy are tested with the Diebold and Mariano (1995) test, with the small-sample correction of Harvey et al (1997). With 21 to 24 forecasts the test has limited power, and we read it with that in mind.

4.5 Drought scenarios

The final step uses the estimated model to simulate the effect of droughts over 2026–2028. The reference scenario assumes normal crop-year rainfall (219 mm) and normal spring temperature in each year. Three drought scenarios change only rainfall: a moderate drought (30 percent below normal in 2027), a severe drought (50 percent below normal in 2027) and a prolonged drought (40 percent below normal in both 2027 and 2028). The last one is close to what Morocco experienced between 2022 and 2024.

For each scenario we compute the loss in agricultural value added relative to the reference path, in percent of annual output. To give an order of magnitude in money, the percentage loss is applied to agricultural value added at current prices in 2023 (154.4 billion dirhams). Uncertainty is measured by drawing the ARDL coefficients 2,000 times from their estimated distribution and reporting the 90 percent interval of the resulting losses. Since the model is log-linear, the relative losses do not depend on the assumption made for 2026.

We did not run a heat-wave scenario. As Section 5.6 shows, the effect of temperature on agricultural value added is not robust at the national quarterly level: its sign depends on which temperature measure is used. A scenario built on such an unstable coefficient would not be credible, and we prefer to report this as a limitation.

5 Results

5.1 Unit root tests

Table 1 summarises the unit root tests. For agricultural value added, the picture depends on the deterministic terms. With a constant only, none of the tests rejects a unit root. Once a trend is added, the ADF, Phillips–Perron and Fourier ADF tests reject at 5 percent, but the KPSS test rejects stationarity at the same level. We read the series as I(1), with the caveat that it is close to trend stationarity. In first differences it is clearly stationary under every test.

The climate variables are stationary in levels. Crop-year rainfall rejects the unit root at 1 percent with all three tests, and so does spring maximum temperature. Mean crop-year temperature is also stationary according to the ADF and PP tests, while the KPSS test with a constant only picks up the warming trend of the last two decades.

The conclusion that matters for what follows is simple: the variables are a mix of I(0) and I(1), and none is I(2). This is the setting for which the ARDL bounds test was designed.

One result deserves comment. The F-statistics on the Fourier terms in the unit root regressions are small, between 1.1 and 6.1, well under the critical values of Enders and Lee (2012). Taken alone, the series do not show the kind of smooth breaks that the Fourier ADF is built to capture. As we show below, the Fourier terms become significant once rainfall is in the model, which suggests that the shifts concern the relation between climate and output rather than either series on its own.

Table 1 Unit root tests, 1998–2025

Variable	Specification	ADF	PP	KPSS	Fourier ADF (k^*)
ln VA (quarterly), level	Constant	−1.06	−0.84	1.60***	−1.85 (5)
	Constant + trend	−3.68**	−3.79**	0.20**	−4.47** (1)
ln VA, first difference	Constant	−7.05***	−12.63***	0.04	−7.82*** (4)
ln rainfall (annual), level	Constant	−4.77***	−4.92***	0.23	−6.15*** (1)
Spring Tmax (annual), level	Constant	−6.90***	−6.74***	0.26	−7.23*** (1)
Crop-year temperature, level	Constant + trend	−5.56***	−5.73***	0.15*	−1.86 (4)

Note: ***, **, * denote rejection at 1, 5 and 10 percent. For ADF, PP and Fourier ADF the null is a unit root; for KPSS the null is stationarity. Fourier ADF critical values from Enders and Lee (2012). Climate variables are tested on 28 annual observations.

Table 2 Fourier ARDL(1, 3, 1), $k = 1$, 1998Q1–2025Q4

	Coefficient	Std. error
Long run		
ln rainfall	0.439***	0.082
Spring Tmax	0.028*	0.015
Short run		
Δ ln rainfall	0.306***	0.034
Δ Spring Tmax	0.023***	0.007
Error-correction term	−0.432***	0.081
$\sin(2\pi t/T)$	−0.035**	0.014
$\cos(2\pi t/T)$	−0.005	0.010
Trend	0.004***	0.001
Bounds test F	7.48	I(0)/I(1) bounds at 1%: 5.19 / 6.34
Adjusted R^2	0.978	$N = 109$

Note: ***, **, * significant at 1, 5 and 10 percent. Long-run standard errors by the delta method. Bounds test, case V of Pesaran et al (2001).

5.2 The Fourier ARDL estimates

The Akaike criterion selects an ARDL(1, 3, 1) with a constant and a trend, whatever the Fourier frequency. Adding Fourier terms with $k = 1$ lowers the AIC from −312.6 to −315.7, and higher frequencies do worse. The two Fourier terms are jointly significant in the ARDL (Wald $F = 3.22$, $p = 0.044$), so we keep them. The sine term is negative and significant at 5 percent; the cosine term is not significant on its own.

The bounds test gives $F = 7.48$, above the upper bound at the 1 percent level (6.34). There is a long-run relationship between agricultural value added and crop-year climate. We note that the presence of Fourier terms changes the distribution of the bounds statistic; given how far the statistic is above the 1 percent bound, the conclusion is unlikely to be affected, but a bootstrap version of the test (Yilanci et al, 2020) would settle the point.

The rainfall coefficients are the core result (Table 2). In the long run, a 10 percent fall in crop-year rainfall lowers agricultural value added by about 4.4 percent. The impact effect is already large: about 3.1 percent in the same year. The error-correction

Table 3 Diagnostic tests on the Fourier ARDL residuals

Test	Statistic	p-value
Breusch–Godfrey LM(1)	0.40	0.529
Breusch–Godfrey LM(4)	0.70	0.952
Breusch–Pagan	13.52	0.196
ARCH(4)	6.48	0.166
CUSUM	0.50	0.965
Jarque–Bera	84.04	0.000
Ramsey RESET	6.02	0.016
BDS, dimension 2 / 3 / 4 / 5	1.67 / 2.48 / 2.60 / 2.68	0.095 / 0.013 / 0.009 / 0.007

term of -0.43 means that a little over 40 percent of any gap between output and its long-run level is closed each quarter, so the economy absorbs most of a rainfall shock within a year. This fits what is known about the sector. Rainfed cereals respond almost entirely within the campaign, while the slower adjustment captures livestock and tree crops, whose recovery takes longer.

The temperature result is harder to read. Spring maximum temperature has a positive coefficient, significant at 1 percent in the short run but only at 10 percent in the long run. A positive effect of heat on output is not what agronomy would lead us to expect. We come back to this in Section 5.6, where we show that the sign is not robust.

5.3 Diagnostics and evidence of nonlinearity

The model passes the usual checks (Table 3). There is no sign of serial correlation at one or four lags, no heteroskedasticity, no ARCH effects, and the CUSUM statistic stays well inside its bounds, so the coefficients are stable over the sample.

Two tests fail, and both are informative. The residuals are not normal (Jarque–Bera $= 84.0$), which is no surprise in a sample that contains several extreme droughts. More interesting is the Ramsey RESET test, which rejects the linear functional form at 5 percent ($p = 0.016$). The BDS test points the same way: independence of the residuals is rejected at dimensions 3, 4 and 5, with p-values below 0.01. Whatever the Fourier ARDL leaves in its residuals is not just noise. It has a nonlinear structure, and this is the statistical basis for adding a neural network to the model.

5.4 Choice of the neural network

On the 2016–2019 validation period, the best configuration turned out to be the simplest one: a single lag of the residuals and three hidden neurons, without the climate variables as additional inputs. It reduced the validation error by about 18 percent relative to the Fourier ARDL alone. Larger networks, and networks fed with climate data, did worse. With around a hundred observations this is what one would expect. A larger network has more weights than the data can pin down, and the climate information is already used by the linear part of the model.

Table 4 Out-of-sample forecast accuracy, 2020Q1–2025Q4

Model	RMSE	MAE	MAPE (%)	DM vs FARDL–ANN
One quarter ahead (24 forecasts)				
ARIMA	2,377	1,676	4.40	1.36
ARIMAX	1,742	1,385	3.68	0.60
ARIMAX–ANN	1,834	1,544	4.09	1.47
ARDL	1,905	1,528	4.12	2.78**
Fourier ARDL	**1,630**	**1,297**	**3.48**	−1.14
Fourier ARDL–ANN	1,668	1,334	3.58	–
Four quarters ahead (21 forecasts)				
ARIMA	4,829	4,216	11.13	1.52
ARIMAX	3,095	2,883	7.68	1.53
ARIMAX–ANN	3,067	2,857	7.61	1.50
ARDL	3,459	2,954	7.92	1.41
Fourier ARDL	2,339	2,061	5.53	1.31
Fourier ARDL–ANN	**2,279**	**2,014**	**5.39**	–

Note: RMSE and MAE in millions of dirhams (volume). DM is the Diebold–Mariano statistic with the Harvey et al (1997) correction; a positive value means the model in that row has larger squared errors than the Fourier ARDL–ANN. ** significant at 5 percent. Best value at each horizon in bold.

5.5 Out-of-sample forecasts

Table 4 reports the accuracy of the six models over 2020–2025, and Figure 3 plots the forecasts of the main models.

Three things come out of the table.

The first is the weight of climate. ARIMA, which ignores rainfall, is the worst model at both horizons. At four quarters its percentage error is twice that of the best model. Everything that follows is secondary to this.

The second is the contribution of the Fourier terms. The Fourier ARDL beats the standard ARDL on every measure and at both horizons, by around 15 percent at one quarter and 30 percent at four. This is also the only difference that the Diebold–Mariano test finds significant. The gradual shifts that the Fourier terms pick up are real, and ignoring them is costly.

The third concerns the neural network, and here the answer depends on the horizon. One quarter ahead, the Fourier ARDL alone is slightly better than the hybrid. Four quarters ahead, the hybrid is the best model on all three measures, with a MAPE of 5.39 percent against 5.53 percent for the Fourier ARDL and 7.68 percent for ARIMAX. The gain is consistent but modest, and it is not significant according to the Diebold–Mariano test, which with 21 forecasts has little power to detect differences of this size. Our reading is that the network adds something at the horizon that matters for budget planning, the coming year, but that the bulk of the improvement over standard models comes from climate and from the Fourier terms. We think it is better to say this plainly than to present the hybrid as a clear winner.

All models miss the jump of 2021Q1 by one quarter. This is not really a forecasting failure. The quarterly series moves in a single step when HCP incorporates a new

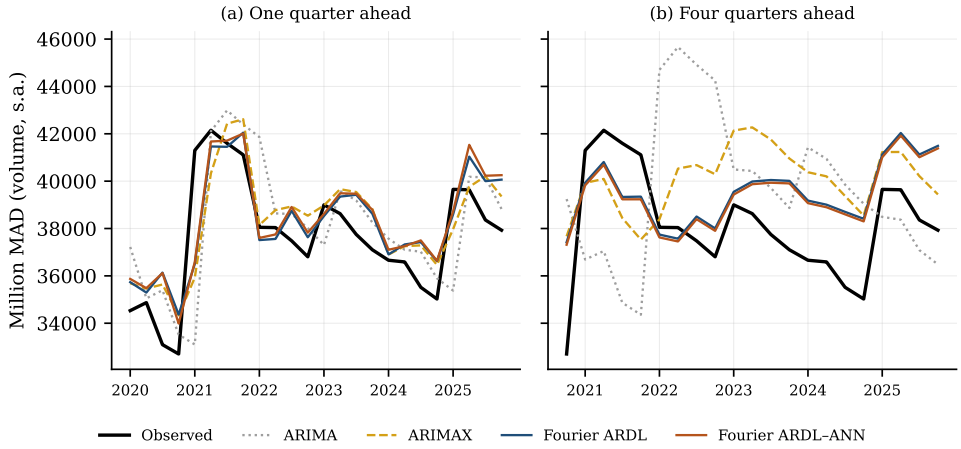


Fig. 3 Out-of-sample forecasts of agricultural value added, 2020–2025
Note: Expanding-window forecasts conditional on observed crop-year climate. ARDL and ARIMAX–ANN omitted for readability.

campaign, and no model that uses the previous quarter's output can anticipate that step without information on the new harvest.

5.6 Robustness

We re-estimated the model with the two alternative temperature measures. The rainfall elasticity barely moves: the long-run coefficient is 0.44 with spring maximum temperature, and 0.36 to 0.37 with mean or maximum temperature over the crop year. Temperature, in contrast, behaves very differently. Its effect is positive with spring maximum temperature but close to zero and insignificant with the other two measures. Rainfall and spring temperature are also negatively correlated across years ($r = -0.35$), so part of what the spring temperature variable captures may simply be the timing of the rainy season.

We also changed the definition of the crop year. With rainfall summed from September to May instead of October to April, the long-run rainfall elasticity is 0.39 instead of 0.44, still significant at 1 percent, and the bounds test still points to cointegration. The fit is worse, however (AIC of -306.0 against -315.7), and the Fourier terms are only significant at 10 percent. September and May rains seem to add more noise than signal at the national level: early autumn rain often falls before sowing, and late spring rain comes too late for most cereals. This is why we keep the October–April window in the main specification.

Our conclusion is that the rainfall effect is solid and that the temperature effect is not identified at this level of aggregation. National averages mix regions and crops whose responses to heat differ in sign, irrigated perimeters among them. Pinning down the effect of heat would require regional or crop-level data.

Table 5 Simulated losses relative to the reference path

Scenario	Year	Rainfall (mm)	Loss in VA (%)	90% interval	Loss (bn MAD)
Moderate drought	2027	154	11.6	[9.5, 13.6]	17.9
(−30% in 2027)	2028	219	2.6	[0.1, 5.2]	4.1
Severe drought	2027	110	21.3	[17.6, 24.7]	32.9
(−50% in 2027)	2028	219	5.0	[0.1, 9.8]	7.8
Prolonged drought	2027	132	16.2	[13.3, 18.9]	25.0
(−40% in 2027 and 2028)	2028	132	19.3	[14.1, 24.1]	29.9

Note: Losses are the percentage gap between annual value added in the scenario and in the reference path (normal rainfall of 219 mm). Intervals come from 2,000 draws of the ARDL coefficients. Money values apply the percentage loss to agricultural value added at current prices in 2023.

5.7 Drought scenarios

Table 5 reports the losses in agricultural value added implied by the hybrid model under the three drought scenarios, relative to a path with normal rainfall, and Figure 4 shows the corresponding paths.

A moderate drought costs around 12 percent of agricultural value added in the year it occurs. This is in line with history: in 2016, rainfall 37 percent below normal went with a 13 percent fall in output, and in 2022, a 31 percent deficit with a 9.5 percent fall. A severe drought roughly doubles the loss, to about 21 percent.

Two further results are worth stressing. The first is persistence. Even when rain returns to normal in 2028, output stays 2.6 to 5 percent below the reference path, because the sector takes time to rebuild herds, plantations and farm finances. The second is that duration matters as much as intensity. A prolonged drought of 40 percent over two years costs about 55 billion dirhams in total, against 41 billion for a single severe year, and the loss in the second year (19.3 percent) is larger than in the first. Morocco went through exactly this kind of sequence between 2022 and 2024. Our results suggest that policies aimed at shortening dry spells in their effects, through water storage, desalination for irrigation and support to herds, may matter more than those aimed at a single bad year.

6 Conclusion and policy implications

This paper set out to forecast Moroccan agricultural value added at the moment when the rain of the crop year is known but its economic effect is not. To do so, we built a quarterly series consistent with the latest national accounts, measured rainfall over the October–April campaign, and combined a Fourier ARDL model with a small neural network fitted to its residuals.

Three results stand out. First, rainfall is by far the main driver of agricultural value added, and its effect is stable across specifications. A 10 percent rainfall deficit lowers output by about 3 percent in the same year and 4.4 percent in the long run. Second, the relation between rainfall and output has shifted gradually over the past twenty-five years, and allowing for these shifts with Fourier terms is what improves forecasts most. Third, a neural network fitted to the residuals improves forecasts a

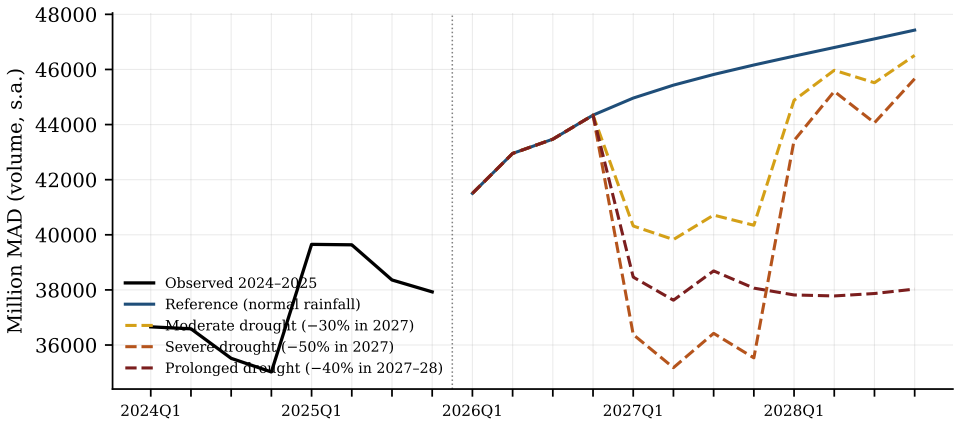


Fig. 4 Agricultural value added under alternative rainfall scenarios, 2026–2028
Note: Paths simulated with the hybrid Fourier ARDL–ANN model. Only crop-year rainfall differs across scenarios; temperature is held at its 1991–2020 norm.

year ahead, where the hybrid model is the most accurate of all the models we tested, but its contribution is small compared with that of climate and of the Fourier terms. We think this is a useful result in itself. In a sample of about a hundred quarters, a well-specified econometric model does most of the work, and the gain from machine learning should be tested rather than assumed.

The scenarios add a further message. Losses from drought do not end when the rain returns: after a single dry year, output remains several percent below its normal path the following year. And a drought that lasts two years at 40 percent below normal rainfall costs more, in total, than a single year at 50 percent below normal.

These results have practical uses. The first concerns forecasting itself. Because the rainfall of a campaign is known by the end of April, a model of this kind can give the Ministry of Economy and Finance, Bank Al-Maghrib and HCP an estimate of agricultural value added for the year several months before the national accounts are available, together with a measure of its uncertainty. It is simple enough to be updated each year and transparent enough to explain to policymakers, which is not always true of more complex machine-learning models.

The second concerns drought policy. Support to farmers in Morocco is usually triggered campaign by campaign, once a drought is declared. Our results suggest that the cost of a drought depends heavily on whether it is followed by another one, and that part of the loss comes after the rains return. This argues for programmes that last beyond the dry year itself, in particular for herds and tree crops, which take longest to rebuild, and for treating water storage, desalination for irrigation and the extension of index-based crop insurance as investments against multi-year droughts rather than responses to single bad years.

The study has several limitations. The quarterly series is partly constructed: HCP spreads each campaign over the four quarters of the year, and we benchmarked the post-2014 quarters to annual totals. The model works at the national level and cannot identify the effect of temperature, whose sign depends on the measure used; regional or

crop-level data would be needed to do so. The test period contains 21 to 24 forecasts, so the Diebold–Mariano test has limited power to distinguish between the best models. Finally, the bounds test with Fourier terms would be best confirmed with bootstrap critical values.

Several extensions follow. Satellite vegetation indices could be added to update forecasts within the season, before the end of April. Regional value added, where it becomes available, would allow the effect of heat and the role of irrigated areas to be measured. Other network architectures, such as recurrent networks, could be tested once longer series exist. And the 2025/26 campaign, one of the wettest in years, will offer a natural test of the model as soon as the climate data for 2026 are published.

Acknowledgements. [To be completed.]

Declarations

- **Funding:** The authors received no specific funding for this work.
- **Competing interests:** The authors declare no competing interests.
- **Data availability:** The quarterly agricultural value added series built for this paper and the climate variables are available from the corresponding author on reasonable request. The underlying data are publicly available from the Haut-Commissariat au Plan and the World Bank Climate Change Knowledge Portal.
- **Code availability:** The code used for estimation, forecasting and simulation is available from the corresponding author on reasonable request.
- **Author contribution:** Lamiaa Aattar: conceptualization, data curation, methodology, formal analysis, writing – original draft. Saad Elouardirhi: supervision, validation, writing – review and editing. Both authors read and approved the final manuscript.

References

Aattar L, Elouardirhi S (2022) The contribution of the agricultural sector to economic growth: a cointegration analysis for Morocco. [Journal to be completed]

Balaghi R, Tychon B, Eerens H, et al (2008) Empirical regression models using NDVI, rainfall and temperature data for the early prediction of wheat grain yields in Morocco. International Journal of Applied Earth Observation and Geoinformation 10(4):438–452

Banerjee P, Arčabić V, Lee H (2017) Fourier ADL cointegration test to approximate smooth breaks with new evidence from crude oil market. Economic Modelling 67:114–124

Becker R, Enders W, Lee J (2006) A stationarity test in the presence of an unknown number of smooth breaks. Journal of Time Series Analysis 27(3):381–409

Bouras EH, Jarlan L, Er-Raki S, et al (2021) Cereal yield forecasting with satellite drought-based indices, weather data and regional climate indices using machine learning in Morocco. Remote Sensing 13(16):3101

Broock WA, Scheinkman JA, Dechert WD, et al (1996) A test for independence based on the correlation dimension. Econometric Reviews 15(3):197–235

Cadenas E, Rivera W (2010) Wind speed forecasting in three different regions of Mexico, using a hybrid ARIMA–ANN model. Renewable Energy 35(12):2732–2738

Daoui M (2023) Macroeconomic forecasting using dynamic factor models: the case of Morocco. Alternatives Managériales et Économiques 5(2):475–492

Denton FT (1971) Adjustment of monthly or quarterly series to annual totals: an approach based on quadratic minimization. Journal of the American Statistical Association 66(333):99–102

Diebold FX, Mariano RS (1995) Comparing predictive accuracy. Journal of Business & Economic Statistics 13(3):253–263

Driouech F, Stafi H, Khouakhi A, et al (2021) Recent observed country-wide climate trends in Morocco. International Journal of Climatology 41(S1):E855–E874

Enders W, Lee J (2012) The flexible Fourier form and Dickey–Fuller type unit root tests. Economics Letters 117(1):196–199

Faruk DÖ (2010) A hybrid neural network and ARIMA model for water quality time series prediction. Engineering Applications of Artificial Intelligence 23(4):586–594

Harvey D, Leybourne S, Newbold P (1997) Testing the equality of prediction mean squared errors. International Journal of Forecasting 13(2):281–291

Hersbach H, et al (2020) The ERA5 global reanalysis. Quarterly Journal of the Royal Meteorological Society 146(730):1999–2049

International Monetary Fund (2026) Leveraging non-traditional data for macroeconomic nowcasting: the case of Morocco. IMF Working Paper 2026/108, International Monetary Fund, Washington, DC

Khashei M, Bijari M (2011) A novel hybridization of artificial neural networks and ARIMA models for time series forecasting. Applied Soft Computing 11(2):2664–2675

Koutroumanidis T, Ioannou K, Arabatzis G (2009) Predicting fuelwood prices in Greece with the use of ARIMA models, artificial neural networks and a hybrid ARIMA–ANN model. Energy Policy 37(9):3627–3634

Makridakis S, Spiliotis E, Assimakopoulos V (2018) Statistical and machine learning forecasting methods: concerns and ways forward. PLoS ONE 13(3):e0194889

McNown R, Sam CY, Goh SK (2018) Bootstrapping the autoregressive distributed lag test for cointegration. Applied Economics 50(13):1509–1521

Pai PF, Lin CS (2005) A hybrid ARIMA and support vector machines model in stock price forecasting. Omega 33(6):497–505

Pesaran MH, Shin Y, Smith RJ (2001) Bounds testing approaches to the analysis of level relationships. Journal of Applied Econometrics 16(3):289–326

Schilling J, Freier KP, Hertig E, et al (2012) Climate change, vulnerability and adaptation in North Africa with focus on Morocco. Agriculture, Ecosystems & Environment 156:12–26

Soumbara S, Moulim A, El Ghini A (2025) Hybrid time series and neural network model for forecasting electricity demand under climate constraints. Journal of Modelling in Management https://doi.org/10.1108/JM2-05-2025-0199

Taskaya-Temizel T, Casey MC (2005) A comparative study of autoregressive neural network hybrids. Neural Networks 18(5–6):781–789

World Bank (2026) World development indicators: NV.AGR.TOTL.ZS and SL.AGR.EMPL.ZS. Accessed October 2026

Yilanci V, Bozoklu S, Gorus MS (2020) Are BRICS countries pollution havens? Evidence from a bootstrap ARDL bounds testing approach with a Fourier function. Sustainable Cities and Society 55:102035

Zhang GP (2003) Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing 50:159–175