sustainability

MDPI

Article

Climatic and Productive Correlates of Coffee Yield in Latin America: Dynamic Adjustment and Long-Run Heterogeneity

Luis Rojas *, Ronald-Kleiner Toledo-Macas and Santiago Ochoa-Moreno

Department of Economics, Universidad Técnica Particular de Loja, Loja 110101, Ecuador; rktoledo@utpl.edu.ec (R.-K.T.-M.); wsochoa@utpl.edu.ec (S.O.-M.)
* Correspondence: larojas@utpl.edu.ec

Abstract

Coffee productivity in Latin America is exposed to climatic variability, input-use pressures, and pronounced cross-country heterogeneity. This study evaluates dynamic associations between coffee yield and climatic, productive, and economic conditions in twelve major producing countries from 1992 to 2023 using an ARDL(1,1) framework with Pooled Mean Group (PMG) and Mean Group (MG) estimators. Both models yield negative and statistically significant adjustment coefficients (−0.866 and −1.087, respectively; $p < 0.001$). Because coffee yield, the constructed NPK-use/coffee-area proxy, and the temperature anomaly are stationary in levels, these parameters are interpreted as dynamic mean reversion within the ECM parameterization rather than as standalone evidence of cointegration; the Westerlund statistics are therefore treated as supplementary diagnostics. A Swamy-type test rejects common long-run slopes (chi-square = 101.62, df = 55, $p < 0.001$), while none of the five long-run covariates is significant at the 5% level in either central model. A targeted denominator-free sensitivity check replacing the area-normalized proxy with national NPK agricultural use in tonnes of product also finds no statistically significant fertilizer association. The results support country-sensitive adaptation and productivity strategies and caution against causal farm-level interpretations of national macro-panel associations.

Keywords: coffee yield; climate variability; dynamic panel; ARDL; long-run heterogeneity; Latin America; agricultural productivity; mean reversion

1. Introduction

Coffee is one of the most important commercial crops in Latin America because of its contribution to rural employment, smallholder income, and agri-food exports. Coffee is a long-lived perennial crop: a coffee field generally requires about three to five years to reach economic yields and can remain productive for 30 years or more, so plantation establishment, renewal, and replacement decisions naturally operate over a multi-year horizon [1]. This biological lag distinguishes coffee from annual crops and is central to the dynamic empirical design used here. At the same time, the sustainability of production is increasingly challenged by climate variability and structural constraints that limit producers' capacity to adjust [2,3].

Despite the growing literature on coffee and climate change, a substantial share of evidence is based on agroecological simulations, spatial suitability analyses, or local case studies. Comparative econometric evidence for Latin America that jointly considers climatic conditions, input intensity, market signals, and dynamic adjustment remains limited [4]. This gap matters because yield responses to changes in climate, fertilization, and economic

Check for updates

Academic Editors: Gabriele Scozzafava and Tommaso Fantechi

Received: 23 July 2026
Revised: 4 September 2026
Accepted: 8 September 2026
Published: 22 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

incentives are unlikely to be instantaneous in a perennial crop. A dynamic framework is therefore needed to distinguish contemporaneous movements from persistent long-run associations while allowing national production systems to differ [5,6].

Building on this gap, the study has three specific objectives. First, it characterizes cross-country and intertemporal variation in coffee yield in twelve Latin American producers over 1992–2023. Second, it estimates short- and long-run associations between yield, a constructed national NPK-use/coffee-area proxy, annual temperature anomaly, the international coffee price, and GDP per capita while explicitly testing whether long-run slopes can be treated as homogeneous. Third, it evaluates robustness to alternative specifications and influential countries. The working hypotheses are that the ARDL-ECM representation exhibits statistically significant dynamic mean reversion (H1) and that long-run slopes are heterogeneous across countries (H2). H1 concerns the estimated adjustment parameter and is not treated as a standalone cointegration test. Because the national panel provides no external source of exogenous variation, all coefficients are interpreted as conditional associations rather than causal effects.

The remainder of the article is organized as follows. Section 2 develops the ecobotanical and economic rationale and reviews the literature. Section 3 describes the data and statistical analysis. Section 4 presents the empirical results. Section 5 discusses the findings and limitations, and Section 6 concludes with the study's novelty and future research priorities.

2. Theoretical Background and Literature Review

2.1. Ecobotanical and Productive Characteristics of Coffea

The genus *Coffea* (*Rubiaceae*) comprises perennial evergreen shrubs or small trees. Coffee fields typically require about three to five years to reach economic yields and may remain in production for 30 years or more [1]. Commercial supply is overwhelmingly dominated by *Coffea arabica* and *Coffea canephora* (*Robusta*), which together account for at least 99.9% of global coffee trade [7]. Arabica is indigenous to the highland forests of Ethiopia and South Sudan [7]. *C. canephora* originated in the humid evergreen forests of central and western sub-Saharan Africa and is now cultivated in about 20 countries; Robusta contributes about 40% of global coffee production, with Arabica accounting for most of the remainder [8]. *Coffea liberica* is indigenous to Upper West Africa (Sierra Leone, Liberia, Côte d'Ivoire, Ghana, and Nigeria) and is cultivated mainly in warm lowland environments, including Malaysia, the Philippines, and Indonesia [9]. *Coffea dewevrei* (excelsa) is indigenous to Central Africa and is cultivated in locations including Guinea, South Sudan, and Uganda [9]. Liberica and excelsa remain extremely minor globally: their combined production is estimated at less than 1000 metric tons, approximately 0.01% of 2024 global coffee exports [9]. Recent genomic evidence supports recognizing *C. dewevrei* as distinct from *C. liberica* [9].

This long biological and productive horizon means that decisions about planting, renovation, fertilization, shade, and other management practices can affect productive capacity with substantial lags. This intertemporal structure motivates a dynamic production framework in which short-run movements need not equal long-run associations.

From the microeconomics of production, coffee yield reflects the interaction between agronomic inputs, technological and management conditions, and environmental constraints. Producers allocate land, labor, capital, and intermediate inputs under uncertain climate and market conditions [5]. Recent work further shows that yield is better understood when agroclimatic and biophysical monitoring variables are considered jointly [10], and process-based evidence indicates nonlinear interactions among temperature, precipitation, radiation, and management [11].

Accordingly, coffee yield is treated here as a productivity outcome rather than as total output. This distinction is important for empirical specification: because FAOSTAT yield is mechanically constructed as production divided by harvested area, harvested area is not included as a separate regressor in the main yield equation. Area dynamics are examined separately in an auxiliary adjustment equation. This removes the direct mechanical dependence that would arise from placing harvested area itself on the right-hand side; however, the constructed NPK-use/coffee-area proxy also uses coffee harvested area as its denominator. That separate shared-denominator concern is addressed explicitly with a denominator-free robustness check based on national NPK agricultural use in tonnes of product.

2.2. Climate Sensitivity and Coffee Productivity

The specialized literature agrees that coffee is among the crops most sensitive to climate change because of its close dependence on specific ranges of temperature, precipitation, humidity, and seasonality. Bunn et al. [12] showed that global warming alters the spatial distribution of suitable areas for *Coffea arabica* and *Coffea canephora* (*Robusta*), shifting cultivation toward higher altitudes and reducing productive stability in several tropical regions. In the same vein, Ovalle-Rivera et al. [13] projected significant losses of climatic suitability in the main producing regions, suggesting that climate pressure on coffee farming is structural rather than merely transitory.

More recently, Grüter et al. [14] showed that coffee suitability depends on the interaction between climatic and edaphic factors, although climatic constraints tend to be more decisive than soil-related factors in the global distribution of the crop. This finding is particularly relevant because it indicates that temperature and precipitation affect not only current observed yields but also the future viability of the production system. Similarly, Byrareddy et al. [15] warn that coffee-related climate risks are sensitive to the spatial scale of analysis, meaning that their impacts cannot be fully understood without considering territorial heterogeneity and differentiated responses across regions and production systems.

Empirical evidence also suggests that the effects of climate on coffee are neither linear nor immediate. Craparo et al. [16] found that sustained increases in minimum night temperatures significantly reduce *Coffea arabica* yields, while Bilen et al. [4], in their systematic review, conclude that climate change simultaneously affects productivity, bean quality, pest incidence, and the stability of ecosystem services associated with coffee farming. This implies that climate effects may accumulate with lags, operate through multiple biological and productive mechanisms, and manifest with varying intensity depending on the time horizon analyzed.

Accordingly, climate change is treated as a structural condition associated with coffee yield. It may alter the production frontier, the efficiency of input use, and producers' adaptive possibilities, but the national-level empirical design does not identify these channels causally. This interpretation provides the conceptual basis for including temperature as a central correlate in the empirical model.

2.3. Productive Inputs, Price Signals, and Area Adjustment

In addition to climate, coffee productivity depends on input-use decisions and on economic conditions that influence investment and management intensity. Price signals need not affect yield contemporaneously: in a perennial crop, they can operate through plantation renewal, tree replacement, area decisions, and capital investment over several years. The international price is retained in the yield equation only as a common market-condition control; the paper does not interpret its coefficient as a structural yield response. Behavioral adjustment to price is examined separately through the harvested-area equation.

Coffee producers combine short-run decisions, such as fertilization and phytosanitary management, with longer-run choices such as replanting and area adjustment. Baffes et al. [2] document strong coffee-price volatility and changes in market organization, but this does not imply an immediate proportional yield response. In the empirical design, price is therefore interpreted cautiously as a common international incentive signal. A four-year lag is used in the auxiliary area equation because it lies within the crop's approximately three-to-five-year establishment horizon [1].

The recent literature also warns that greater input intensity does not always translate into sustained yield increases. Chéron-Bessou et al. [17] show that fertilizers constitute one of the main sources of environmental impact in the agricultural phase of coffee production. In the present national panel, however, an aggregate fertilizer coefficient cannot by itself identify technical efficiency, over-intensification, soil degradation, or management quality; those mechanisms require more detailed agronomic data.

Thus, coffee yield is modeled as a dynamic function of fertilizer intensity, climate, the common international coffee price, and a broad structural-development control. GDP per capita is included as a broad macroeconomic control associated with productive capacity and secular development conditions, not as a direct measure of technological progress. No additional deterministic time trend is included in the main equation, which keeps the specification parsimonious and avoids adding a second broad trending proxy.

2.4. Dynamic Adjustment and Estimator Choice

The empirical model follows the partial-adjustment logic of dynamic heterogeneous panels [6]. Coffee production cannot adjust instantaneously to changes in climate, input use, or economic incentives because biological lags and adjustment costs make productive responses persistent. A one-lag ECM parameterization is therefore used to distinguish contemporaneous changes from the conditional long-run component of the ARDL model. In this study, a negative adjustment coefficient is interpreted as dynamic mean reversion or adjustment in the conditional yield equation; it is not treated by itself as proof that all variables form a cointegrating relationship.

Bilen et al. [4] emphasize that the effects of climate change on coffee farming tend to be cumulative, delayed, and heterogeneous, reinforcing the need for models capable of distinguishing short-run relationships from persistent longer-run associations. In the case of coffee, this distinction is particularly relevant because the impact of a climate anomaly or a change in input use may not be fully reflected in the harvest of the same year, but rather in subsequent periods due to the crop's phenological cycle. Hence, a static model may be insufficient to capture the relationship among climate, inputs, and yield.

Two estimators are emphasized. Pooled Mean Group (PMG) restricts long-run slopes to be common while allowing short-run coefficients and adjustment speeds to differ by country [6]. Mean Group (MG) estimates each country separately and averages the coefficients, imposing no common long-run slope. Standard MG does not include cross-sectional averages; the empirical implementation therefore contains no cross-sectional-average terms. Dynamic Fixed Effects (DFE) is retained only as a robustness benchmark.

Potential endogeneity remains relevant because fertilizer decisions and yield may be jointly determined and because economic conditions can respond to agricultural performance. The dynamic specification establishes temporal ordering but does not generate exogenous variation. Accordingly, the study does not use causal language, and robustness exercises are interpreted as sensitivity analyses rather than as instruments for causal identification.

3. Materials and Methods

3.1. Research Design, Data Construction, and Variables

This study adopts a quantitative, longitudinal, and comparative design based on a balanced country-year panel for twelve major coffee-producing countries in Latin America over 1992–2023: Brazil, Colombia, Peru, Honduras, Guatemala, Mexico, Nicaragua, El Salvador, Costa Rica, Ecuador, Bolivia, and the Dominican Republic. The unit of analysis is the country-year (384 observations before lagging). The national scale is appropriate for regional dynamic comparison but cannot be interpreted as farm-level or subnational evidence.

The empirical dataset combines FAOSTAT, the World Bank Climate Change Knowledge Portal/CRU climate information, World Development Indicators, and the IMF Primary Commodity Prices series distributed through FRED. Coffee yield is reported in metric tons per hectare (t/ha) for presentation. Estimation uses the natural logarithm of the corrected yield series; expressing the same series in kg/ha or t/ha changes only the intercept, not the slope or adjustment coefficients. The annual temperature variable is a country-level national anomaly in degrees Celsius; it is not a separate anomaly calculated for each coffee-growing region. This limitation is stated explicitly because national climate averages may smooth conditions in high-altitude or otherwise specialized coffee zones.

The final explanatory variables are a constructed NPK-use/coffee-area proxy, annual temperature anomaly and its squared term, the international price of Other Mild Arabica coffee, and GDP per capita in constant 2015 US dollars. The fertilizer numerator was independently verified against the official FAOSTAT Fertilizers by Product metadata. The underlying item is 'NPK fertilizers', a compound-fertilizer product category (HS 310520), and the relevant element is Agricultural Use, reported in tonnes of product [18]. FAOSTAT defines agricultural use at national level as use for crops, livestock, forestry, fisheries, and aquaculture, excluding animal feed; because some other plant-nutrition uses cannot always be separated, they may also be included [18]. The numerator is therefore national agricultural use and is not coffee-specific. In the supplied panel, the constructed proxy equals national NPK agricultural use (tonnes of product) divided by coffee harvested area (ha), so its unit is tonnes of NPK product per hectare of coffee harvested area. It is interpreted only as a scale-normalized national input proxy, not as the actual NPK application rate on coffee farms. The official WDI series NY.GDP.PCAP.KD is used for GDP per capita (Table 1). Harvested area is retained for descriptive and auxiliary area-adjustment analysis but excluded as a separate regressor from the main yield equation. A denominator-free robustness specification uses ln(1 + national NPK agricultural use in tonnes of product), without division by coffee harvested area. Pesticides, precipitation, irrigation, producer prices, and agricultural employment are not estimated in the final specification. Econometric estimation and data processing were conducted using Stata 18 (StataCorp LLC, College Station, TX, USA; https://www.stata.com). No specialized laboratory equipment was used.

3.2. Dynamic Empirical Specification

The main yield equation is deliberately parsimonious. Harvested area is excluded as a separate regressor to remove the direct mechanical link created by the construction of the yield variable. This step does not by itself resolve the distinct shared-denominator issue created by the constructed NPK-use/coffee-area proxy; Section 3.4 therefore introduces a denominator-free sensitivity check using national NPK agricultural use in tonnes of product, without division by coffee harvested area. Unrestricted year fixed effects are also excluded from the price-containing specification because the international coffee price has

exactly the same value for all countries in a given year; including both would make price perfectly collinear with the full set of year effects. Country intercepts are retained.

$$\ln(Y_{it}) = \alpha_i + \beta_1 \ln(1 + F_{it}) + \beta_2 T_{it} + \beta_3 T_{it}^2 + \beta_4 \ln(Price_t) + \beta_5 \ln(GDPpc_{it}) + \varepsilon_{it} \quad (1)$$

where Y_{it} is coffee yield; F_{it} is the constructed NPK-use/coffee-area proxy, defined as national NPK agricultural use in tonnes of product divided by coffee harvested area in hectares; T_{it} is the country-level annual temperature anomaly; $Price_t$ is the common international Other Mild Arabica price; $GDPpc_{it}$ is GDP per capita; $alpha_i$ is a country-specific intercept; and $epsilon_{it}$ is the idiosyncratic error. Because the NPK numerator covers national agricultural use rather than coffee-specific use [18], F_{it} is interpreted as an aggregate scaling proxy, not as a measured coffee-farm fertilizer dose. No unrestricted year effect and no deterministic time trend are included in Equation (1).

Table 1. Variables, sources, transformations, and empirical role.

Variable	Definition/Measurement	Source	Transformation	Empirical Role
Coffee yield	Green coffee yield (t/ha)	FAOSTAT	ln(corrected yield)	Dependent variable
Constructed NPK-use/coffee-area proxy	National agricultural use of NPK fertilizers (tonnes of product)/coffee harvested area (ha); unit = t product per ha coffee area; numerator is not coffee-specific [18]	FAOSTAT Fertilizers by Product (NPK fertilizers, Agricultural Use) [18]; authors' calculation	ln(1 + x)	Scale-normalized national input proxy
National NPK agricultural use (robustness only)	FAOSTAT NPK fertilizers, agricultural use; tonnes of product; national agricultural coverage; no coffee-area denominator [18]	FAOSTAT Fertilizers by Product [18]	ln(1 + x)	Denominator-free sensitivity check
Temperature anomaly	Annual national temperature anomaly (degrees C)	World Bank CCKP/CRU	Level and squared term	Climate correlate
International coffee price	Other Mild Arabica price (US cents/lb)	IMF via FRED	ln	Common market control
GDP per capita	GDP per capita (constant 2015 US$)	WDI, NY.GDP.PCAP.KD	ln	Structural-development control
Harvested area	Coffee harvested area (ha)	FAOSTAT	ln	Descriptive and auxiliary area equation only

Note. CCKP = Climate Change Knowledge Portal; CRU = Climatic Research Unit; IMF = International Monetary Fund; WDI = World Development Indicators. The temperature anomaly is national for each country-year, not coffee-region specific. Harvested area is not included as a separate regressor in the main yield equation. The fertilizer numerator is FAOSTAT Fertilizers by Product, item 'NPK fertilizers' (HS 310520), element Agricultural Use, measured in tonnes of product [18]. Its coverage is national agricultural use rather than coffee-specific use. The constructed proxy divides that numerator by coffee harvested area (ha), yielding tonnes of NPK product per hectare of coffee harvested area; it is a scaling proxy and not a measured coffee-farm application rate.

The nonlinear temperature term allows the conditional association to vary with the anomaly. Because both temperature coefficients prove statistically imprecise, the analysis does not treat the implied curve as a substantive climate-response function.

$$\partial \ln(Y)/\partial T = \beta_2 + 2\beta_3 T \quad (2)$$

The estimated dynamic specification is an ARDL(1,1) written in error-correction form. On a common 1995–2023 comparison sample, the Bayesian Information Criterion (BIC) is 928.93 for ARDL(1,1), lower than all alternatives considered with p and q up to 3; the next-best BIC is 932.59 for ARDL(2,1). The Akaike criterion differs by only 0.20 between

those two models. Given T = 32 and the need to avoid country-level over-parameterization, the BIC-preferred ARDL(1,1) is used throughout.

$$\Delta y_{it} = \varphi_i[y_{i,t-1} - \theta_i' X_{i,t-1}] + \gamma_i' \Delta X_{it} + \mu_i + e_{it} \quad (3)$$

Here $y_{it} = \ln(Y_{it})$ and X_{it} contains $\ln(1 + F_{it})$, T_{it}, T_{it2}, $\ln(Price_t)$, and $\ln(GDPpc_{it})$. Under PMG, theta is restricted to be common across countries while phi_i and $gamma_i$ remain heterogeneous; under MG, $theta_i$, phi_i, and $gamma_i$ are all country-specific. The reported short-run coefficients are the cross-country means of $gamma_i$.

$$\ln(A_{it}) = c_i + \rho \ln(A_{i,t-4}) + \kappa \ln(Price_{t-4}) + u_{it} \quad (4)$$

Equation (4) is an auxiliary Nerlove-type harvested-area adjustment equation. It is not part of the yield regression. The four-year lag reflects the crop's establishment horizon and allows price signals to affect planting/area decisions separately from yield. This separation directly avoids interpreting area expansion as a productivity determinant when yield already contains harvested area in its denominator.

3.3. Statistical Analysis, Diagnostics, and Model Selection

The central estimates use PMG and MG. PMG is efficient only if the common long-run slope restriction is acceptable; MG is less restrictive and permits country-specific long-run coefficients [6]. DFE is used as a robustness benchmark, not as a third co-equal central model. The homogeneity restriction is evaluated with a Swamy-type Wald test [19]. It assesses cross-country slope homogeneity and is conceptually distinct from a Hausman specification comparison.

Model choice is therefore based on the slope-homogeneity evidence rather than on a Hausman specification comparison. If the Swamy-type null of common long-run slopes is rejected, MG receives greater interpretive weight, while PMG remains a restricted benchmark. Significance is reported explicitly for every long-run and short-run coefficient (Table 2).

$$H_0: \theta_i = \theta \text{ for all countries (long-run slope homogeneity)} \quad (5)$$

Table 2. Dynamic estimators and interpretation in the empirical strategy.

Estimator	Long-Run Slopes	Short-Run Dynamics	Main Role	Interpretation
PMG	Common	Country-specific	Restricted benchmark	Use if homogeneity is tenable
MG	Country-specific	Country-specific	Main heterogeneity estimator	Preferred when homogeneity is rejected
DFE	Common	Restricted	Robustness only	Sensitivity benchmark

Note. PMG = Pooled Mean Group; MG = Mean Group; DFE = Dynamic Fixed Effects. Standard MG does not contain cross-sectional averages.

Before dynamic estimation, stochastic properties are evaluated with the cross-sectionally augmented IPS (CIPS) procedure based on CADF regressions for the country-varying variables [20]. Because the international coffee price is a single common time series replicated across countries, it is tested separately with univariate Augmented Dickey–Fuller (ADF) and KPSS procedures [21,22]. This directly evaluates whether the price series is at most I(1), which is necessary for the ARDL framework.

$$CIPS = \frac{1}{N} \Sigma_{i=1}^{N} CADF_i \quad (6)$$

$$\Delta y_{it} = a_i + b_i y_{i,t-1} + c_i \bar{y}_{t-1} + d_i \Delta \bar{y}_t + \Sigma_{j=1}^{p} \gamma_{ij} \Delta y_{i,t-j} + \varepsilon_{it} \quad (7)$$

The ARDL framework allows a mixture of I(0) and I(1) variables but excludes I(2) processes. This admissibility condition is distinct from establishing cointegration among the variables. The main price series does not reject the unit-root null in levels under the ADF test, but its first difference is strongly stationary; the complementary KPSS test rejects level stationarity and does not reject stationarity in first differences. Taken together, the univariate evidence supports treating the price series as I(1), not I(2).

Cross-sectional dependence is examined with Pesaran's CD test [23] after estimation. The standard four-statistic Westerlund error-correction panel procedure (Gt, Ga, Pt, Pa) is retained as a supplementary diagnostic with BIC-selected auxiliary lags from 0 to 1 [24]. However, because the integration profile is mixed and coffee yield, fertilizer intensity, and temperature are I(0), the Westerlund statistics are not treated as establishing cointegration or as validating H1. Their nominal null is no error correction, and the four statistics are reported directly with asymptotic left-tail inference; no bootstrap *p*-values are claimed.

$$CD = \left(\frac{2T}{N(N-1)}\right)^{1/2} \Sigma_{i=1}^{N-1} \Sigma_{j=i+1}^{N} \hat{\rho}_{ij} \quad (8)$$

The standard Westerlund statistics are reported jointly rather than selectively (Table 3). In the estimation sample, three of four asymptotic statistics reject their nominal null, while Ga does not. Given the mixed I(0)/I(1) integration profile, these outcomes are treated only as supplementary diagnostics and are not described as proof of a cointegrating relation. H1 is evaluated from the sign and significance of the model-specific adjustment coefficients, which are interpreted as dynamic mean reversion in the conditional ARDL equation. Residual CD tests are then used to verify that common residual dependence is not statistically detectable in the reported PMG, MG, and DFE specifications.

Table 3. Diagnostic sequence and decision criteria.

Diagnostic	Null/Criterion	Decision Rule	Purpose	Reported Result Type
Cross-sectionally augmented unit-root test (CIPS/CADF)	Unit root	Reject at level or first difference	Country-varying integration order	Panel statistic
Augmented Dickey–Fuller (ADF), price	Unit root	First difference must reject	Common price integration order	Univariate statistic
Kwiatkowski–Phillips–Schmidt–Shin (KPSS), price	Stationarity	First difference should not reject	Cross-check price order	Univariate statistic
Westerlund	No error correction (standard Westerlund framework)	Report nominal result only; do not use to validate H1 under mixed I(0)/I(1)	Supplementary sensitivity diagnostic	Gt, Ga, Pt, Pa
Pesaran cross-sectional dependence (CD)	Cross-sectional independence	Non-rejection preferred in residuals	Residual common dependence	CD statistic and *p*-value
Swamy-type Wald	Common long-run slopes	Reject favors heterogeneous interpretation	Estimator choice	Chi-square and *p*-value

Note. CIPS = cross-sectionally augmented IPS; CADF = cross-sectionally augmented Dickey–Fuller; ADF = Augmented Dickey–Fuller; KPSS = Kwiatkowski–Phillips–Schmidt–Shin; CD = cross-sectional dependence. Test definitions and implementation are detailed in Section 3.3.

3.4. Robustness and Methodological Scope

Robustness analysis focuses on the main identification and specification issues of the empirical design. First, DFE versions of the no-area yield model are estimated after excluding each country in turn. Second, a price-excluded specification includes unrestricted year fixed effects; price must be omitted in that exercise because it is common within each

year. Third, the auxiliary four-year area equation separates planting/area adjustment from the yield equation. Fourth, to isolate the shared-denominator concern raised by the constructed NPK-use/coffee-area proxy, a targeted DFE-style ARDL(1,1) sensitivity check replaces that proxy with the denominator-free FAOSTAT numerator: national Agricultural Use of NPK fertilizers, measured in tonnes of product [18], transformed as ln(1 + NPK). This auxiliary check estimates dynamic yield adjustment with lagged yield, lagged NPK, lagged temperature and temperature squared, their contemporaneous differences, country effects, and unrestricted year effects, with inference clustered by country. The common international price is absorbed by the unrestricted year effects, and GDP per capita is not included in this targeted regression. The check is designed specifically to isolate denominator sensitivity and does not replace the central PMG-MG specification.

For logarithmic regressors, long-run coefficients are elasticities of association; temperature coefficients are semi-elasticities because the anomaly remains in degrees Celsius. No coefficient on harvested area is interpreted in the main yield model because harvested area has been removed as a separate regressor. This removes the direct area regressor but does not eliminate the possible shared-denominator dependence between yield and the constructed NPK-use/coffee-area proxy. The denominator-free NPK sensitivity check is therefore reported separately and uses national NPK agricultural use in tonnes of product without the coffee-area denominator.

Endogeneity remains a limitation because fertilizer intensity, macroeconomic conditions, and yield can be jointly determined. Lagging and the ECM parameterization organize temporal dynamics but do not establish exogeneity or cointegration. Accordingly, the analysis estimates conditional associations and avoids causal language.

The empirical results should be interpreted within the limits of a national macro-panel. The model cannot capture subnational variation in altitude, shade systems, soil quality, plantation age, varieties, irrigation technology, farm size, or agronomic practices. The national temperature anomaly may differ from conditions in specific coffee-growing zones. The empirical design addresses key specification threats by removing harvested area as a separate yield regressor, testing the common price separately, treating Westerlund statistics only as supplementary under the mixed integration profile, distinguishing PMG from MG, reporting short-run coefficients, adding a denominator-free NPK sensitivity check, and checking residual cross-sectional dependence (Table 4).

Table 4. Robustness checks and analytical purpose.

Robustness Check	Purpose	Interpretive Value
Leave-one-country-out Dynamic Fixed Effects (DFE)	Assess influential national cases	Tests coefficient sensitivity
Year fixed effects without common price	Absorb unrestricted common time shocks	Avoids price/time-FE collinearity
Four-year area-adjustment equation	Separate planting response from yield	Avoids denominator dependence
National NPK agricultural use, DFE-style ARDL(1,1)	Assess shared-denominator sensitivity	Uses ln(1 + NPK agricultural use in tonnes of product) without division by coffee harvested area [18]

Note. DFE = Dynamic Fixed Effects. These exercises are robustness checks and do not replace the central PMG-MG comparison.

4. Results

This section reports results from the specified yield model and the audited panel. All central estimates exclude harvested area from the yield equation, use ARDL(1,1), and distinguish the restricted PMG benchmark from the heterogeneous MG estimator. DFE appears only in robustness diagnostics.

4.1. Data Audit, Corrected Sample, and Descriptive Structure

Table 5 summarizes the final audited variables without reproducing lengthy unit-conversion steps. Yield is displayed in t/ha, and GDP per capita is taken directly from the official WDI series. The temperature anomaly is expressed in degrees Celsius. All reported descriptive statistics and coefficient tables use consistent rounding.

Table 5. Data audit and correction decisions.

Variable Audited	N	Minimum	Maximum	Final Decision
Coffee yield	384	0.01 t/ha	1.77 t/ha	Use t/ha for presentation; ln in estimation
Temperature anomaly	384	−0.10 C	2.19 C	Use national annual anomaly
GDP per capita	384	595	14,532	Official WDI NY.GDP.PCAP.KD
Other Mild Arabica price	384	61.46	271.11	Complete annual common series

Note. WDI = World Development Indicators. Values are rounded for presentation. GDP per capita is in constant 2015 US dollars; the coffee price is in US cents per pound. The estimation results are not presented with excessive decimal precision.

Table 6 shows substantial cross-country and intertemporal dispersion. Mean yield is 0.66 t/ha (SD = 0.39), while harvested area varies from 22,136 to 2,500,320 ha. The constructed NPK-use/coffee-area proxy is strongly dispersed relative to its mean, and the national annual temperature anomaly ranges from −0.10 to 2.19 degrees C. This heterogeneity supports a dynamic panel design that allows national adjustment patterns to differ.

Table 6. Descriptive statistics after cleaning.

Variable	N	Mean	SD	Min	Max
Yield (t/ha)	384	0.66	0.39	0.01	1.77
Harvested area (ha)	384	429,980.64	562,882.83	22,136	2,500,320
Constructed NPK-use/coffee-area proxy (t product/ha coffee area)	384	0.08	0.23	0.00	2.31
Temperature anomaly (C)	384	0.74	0.43	−0.10	2.19
Coffee price (US cents/lb)	384	141.29	53.17	61.46	271.11
GDP per capita (2015 US$)	384	5019.29	2867.75	594.69	14,531.50

Note. SD = standard deviation. Values are rounded to two decimals where meaningful. Harvested area is described but is not a regressor in the main yield model.

Figure 1 shows pronounced but temporary regional movements rather than a monotonic path. The unweighted mean falls by 18.4% in 1996 relative to 1995, by 17.8% in 2012 relative to 2011, and by 18.7% in 2020 relative to 2019; it then rises by 10.1% in 2021 (remaining below its 2019 level) and by 30.6% in 2022. Inspection of the underlying country-year panel shows that these regional dips arise from simultaneous sharp observations in several countries rather than from a single region-wide value. Because the study is observational and annual, the figure does not identify a unique causal event for each decline; the analysis therefore reports the data pattern instead of attributing those years to unsupported historical causes.

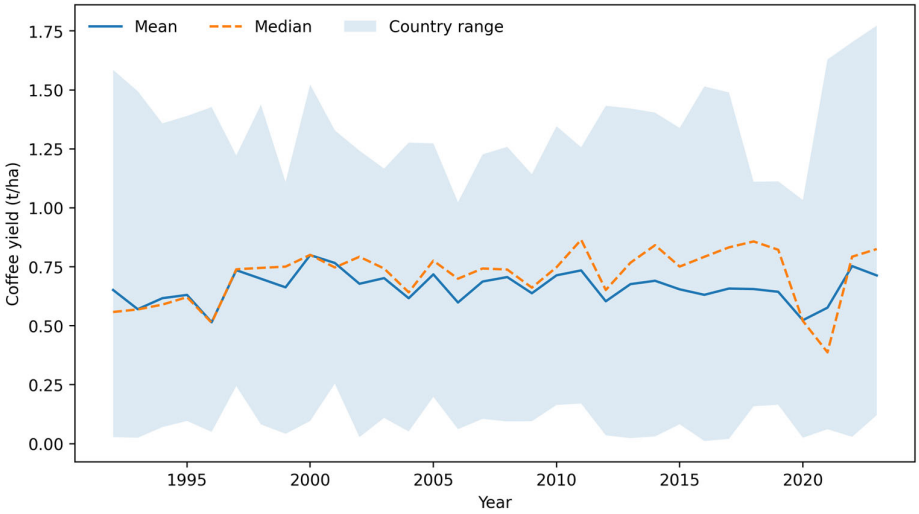

Figure 1. Evolution of coffee yield in Latin America, 1992–2023. The solid line is the unweighted country mean, the dashed line is the median, and the shaded band is the observed country range. Yield is expressed in t/ha.

Figure 2 confirms that national productivity regimes differ markedly in level and timing. The persistence of these differences motivates explicit testing of the long-run homogeneity restriction rather than assuming a single regional slope a priori.

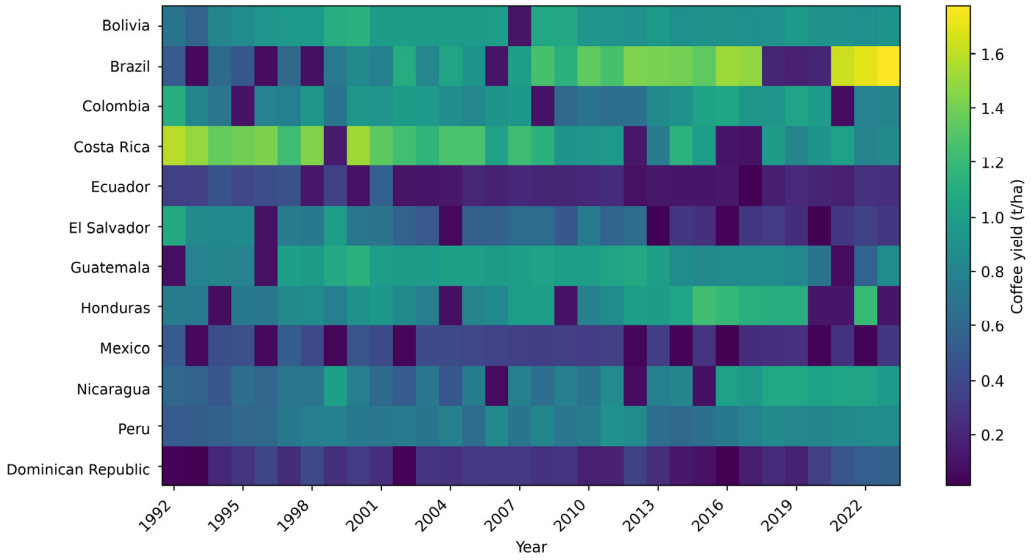

Figure 2. Country-year heterogeneity in coffee yield (t/ha), 1992–2023. Each cell is a country-year observation from the audited panel.

4.2. Integration Properties and Cross-Sectional Diagnostics

The panel diagnostics in Table 7 indicate that the country-varying variables are I(0) or I(1): coffee yield, the constructed NPK-use/coffee-area proxy, and temperature are stationary in levels, whereas GDP per capita is treated as I(1). Because the temperature anomaly is stationary in levels, its squared transformation is also treated as stationary under the maintained finite-moment conditions. The common international price is treated separately rather than replicated across countries for panel unit-root inference; its univariate tests are reported in Table 8. These integration results also motivate the qualified interpretation of H1: the negative adjustment coefficient is treated as dynamic mean reversion in the conditional ARDL-ECM and not as independent proof of cointegration.

Table 7. CIPS panel unit-root tests based on CADF regressions.

Variable	CIPS Level	Level Decision	CIPS First Difference	First-Difference Decision	ARDL Implication
ln_yield	−3.5832	Reject unit root	−5.9906	Reject unit root	I(0) admissible
ln1p_fert	−2.6691	Reject unit root	−5.0051	Reject unit root	I(0) admissible
temp anomaly	−4.2341	Reject unit root	−6.3947	Reject unit root	I(0) admissible
ln_gdppc	−2.2885	Do not reject	−3.4192	Reject unit root	I(1) admissible

Note. CIPS = cross-sectionally augmented IPS; CADF = cross-sectionally augmented Dickey–Fuller; IPS = Im–Pesaran–Shin. The CIPS procedure follows Pesaran [20]. The common international price is intentionally excluded from this panel test and tested as a single time series in Table 8.

Table 8. Univariate price integration and supplementary Westerlund diagnostics.

Diagnostic	Statistic	*p*-Value	Decision	Interpretation	Specification
ADF price, level	−2.218	0.200	Do not reject unit root	Nonstationary in level	Intercept
ADF price, first difference	−4.965	<0.001	Reject unit root	Stationary first difference	Intercept
KPSS price, level	0.483	0.046	Reject stationarity	Consistent with I(1)	Level stationarity null
KPSS price, first difference	0.071	>0.100	Do not reject stationarity	No evidence of I(2)	Stationarity null
Westerlund Gt	−4.244	<0.001	Reject nominal null	Supplementary only; mixed-order panel	BIC lags 0–1
Westerlund Ga	−14.791	0.531	Do not reject nominal null	Supplementary only; mixed-order panel	BIC lags 0–1
Westerlund Pt	−17.068	<0.001	Reject nominal null	Supplementary only; mixed-order panel	BIC lags 0–1
Westerlund Pa	−17.872	0.003	Reject nominal null	Supplementary only; mixed-order panel	BIC lags 0–1

Note. ADF = Augmented Dickey–Fuller; KPSS = Kwiatkowski–Phillips–Schmidt–Shin; BIC = Bayesian Information Criterion. Westerlund statistics are the standard error-correction tests of [24]; reported *p*-values are asymptotic left-tail values, not bootstrap *p*-values. Auxiliary lags 0–1 are selected by BIC, with an average selected lag of 0.33. Because the panel contains a mixture of I(0) and I(1) variables, these statistics are reported as supplementary diagnostics and are not used to establish cointegration or validate H1.

After estimation, residual Pesaran CD statistics are −0.848 ($p = 0.396$) for PMG, 0.125 ($p = 0.900$) for MG, and −0.550 ($p = 0.582$) for DFE. Thus, residual cross-sectional dependence is not statistically detected in any of the three reported specifications.

4.3. Price Integration, Supplementary Westerlund Diagnostics, and Dynamic Adjustment

Table 8 reports two diagnostic blocks. First, the common price series is stationary after first differencing under ADF and is consistent with I(1) under the complementary KPSS evidence; there is therefore no evidence that it is I(2). Second, the standard four-statistic Westerlund procedure is reported directly [24]. Three of four statistics reject their nominal null, whereas Ga does not. Because the panel mixes I(0) and I(1) variables and the dependent variable is stationary in levels, these rejections are not interpreted as establishing a cointegrating relationship. The model-specific adjustment coefficients are therefore discussed as dynamic mean-reversion parameters in the ECM representation, not as independent cointegration tests.

4.4. Dynamic Panel Estimates: PMG and MG

Table 9 shows negative and highly significant mean adjustment coefficients in both central specifications. PMG (−0.866) lies between −1 and 0, consistent with monotonic

annual mean reversion in the conditional dynamic equation. The MG mean (−1.087) lies just below −1, consistent with mild overcorrection and oscillatory adjustment. Because coffee yield is I(0), these coefficients are not interpreted as proof of convergence to a cointegrating equilibrium. They are statistical country-year adjustment parameters, not biological rates for individual coffee trees.

Table 9. Dynamic panel model summary.

Model	Observations	Countries	Mean Adjustment Coefficient (ECM Form)	SE	Statistic	*p*-Value	95% CI Lower	95% CI Upper
PMG	372	12	−0.8657	0.1011	−8.56	<0.001	−1.0639	−0.6675
MG	372	12	−1.0873	0.0910	−11.95	<0.001	−1.2656	−0.9090

Note. ECM = error-correction-model parameterization; PMG = Pooled Mean Group; MG = Mean Group; SE = standard error. The reported ECM term is interpreted as a dynamic-adjustment coefficient and not, by itself, as evidence of cointegration. PMG and MG are the two central estimators. Confidence intervals use coefficient +/− 1.96 SE for compact reporting.

Table 10 contains the two central estimators, PMG and MG. No long-run covariate is statistically significant at the 5% level in either model. The temperature linear term is positive and the squared term negative in both specifications, but neither is statistically precise; consequently, no nonlinear temperature–response curve is interpreted as a substantive finding. The inference therefore centers on dynamic adjustment and heterogeneity rather than on an unstable individual slope.

Table 10. Long-run coefficients across dynamic estimators.

Variable	PMG LR Coef.	PMG SE	PMG *p*-Value	MG Mean LR Coef.	MG SE	MG *p*-Value	Significant at 5%?
ln(1 + NPK-use/coffee-area proxy)	0.0528	0.3563	0.882	0.8172	0.7641	0.308	No
Temperature anomaly	0.0789	0.5171	0.879	0.1209	0.5870	0.841	No
Temperature anomaly squared	−0.4561	0.3038	0.133	−0.3874	0.4237	0.380	No
ln(coffee price)	0.0769	0.1554	0.621	0.1627	0.1833	0.394	No
ln(GDP per capita)	0.1401	0.3606	0.698	0.4657	0.6891	0.513	No

Note. LR = long run; PMG = Pooled Mean Group; MG = Mean Group; SE = standard error. The temperature variable is the national annual anomaly, not absolute temperature. The final column indicates whether the variable is significant at 5% in either central model.

Fertilizer intensity, price, and GDP per capita likewise fail to show a stable statistically significant long-run association across PMG and MG. This does not establish a zero agronomic or economic effect; it indicates that the national macro-panel cannot identify a common regional slope with sufficient precision.

The climate result is interpreted with particular caution because annual national anomalies aggregate heterogeneous elevations, microclimates, shade systems, and phenological conditions. The variable labeled temperature in all tables is explicitly the annual temperature anomaly in degrees Celsius.

The Swamy-type statistic is 101.62 with 55 degrees of freedom ($p = 0.00013$). The null of homogeneous long-run slopes is rejected. Accordingly, MG receives greater interpretive weight, while PMG is retained as a restricted benchmark (Table 11).

Table 11. Swamy-type Wald test of long-run slope homogeneity.

Test	Chi-Square	df	*p*-Value	Decision
Swamy-type long-run slope homogeneity	101.6218	55	0.00013	Reject common long-run slopes

Note. df = degrees of freedom. The null hypothesis is common long-run slopes across countries. Inference is based on the Swamy-type Wald statistic [19].

4.5. Short-Run Dynamics and Robustness

Table 12 reports the short-run coefficients explicitly. None of the mean short-run coefficients is statistically significant at 5%, reinforcing the conclusion that the strongest empirical signal is dynamic adjustment or mean reversion rather than a common contemporaneous input or climate response.

Table 12. Mean short-run coefficients from the ARDL(1,1) models.

Variable	PMG SR Mean	PMG SE	PMG *p*-Value	MG SR Mean	MG SE	MG *p*-Value
Delta ln(1 + NPK-use/coffee-area proxy)	1.3871	0.9836	0.186	1.7662	1.1576	0.155
Delta temperature	0.1951	0.2741	0.491	0.2222	0.4781	0.651
Delta temperature squared	−0.2099	0.1872	0.286	−0.2249	0.3335	0.514
Delta ln(coffee price)	0.2071	0.2497	0.424	0.1878	0.2546	0.476
Delta ln(GDP per capita)	−1.5189	1.6865	0.387	−1.8361	2.1997	0.422

Note. SR = short run; PMG = Pooled Mean Group; MG = Mean Group; SE = standard error; ARDL = Autoregressive Distributed Lag. These are cross-country mean contemporaneous-difference coefficients from the ARDL(1,1) error-correction representation. SE and *p*-values summarize cross-country dispersion of the country-specific short-run estimates; they are not pooled information-matrix standard errors. None is statistically significant at 5%.

Robustness checks support the central interpretation. In a price-excluded DFE specification with unrestricted year fixed effects, the adjustment coefficient in the ECM parameterization remains −0.931 ($p < 0.001$), while the long-run fertilizer, temperature, temperature-squared, and GDP-per capita slopes remain statistically insignificant. Thus, the main dynamic-adjustment finding is not dependent on incorrectly combining a common price series with unrestricted year dummies.

To address the shared-denominator concern directly, a targeted denominator-free DFE-style ARDL(1,1) check replaces the constructed NPK-use/coffee-area proxy with the verified FAOSTAT numerator: national Agricultural Use of NPK fertilizers, measured in tonnes of product [18], transformed as ln(1 + NPK). The specification uses 372 observations, country effects, unrestricted year effects, temperature and temperature-squared terms, and country-clustered inference. The adjustment coefficient is −0.9266 ($p < 0.001$). The implied long-run coefficient on ln(1 + NPK) is −0.0192 (SE = 0.0143, $p = 0.207$), and the short-run coefficient on the first difference of ln(1 + NPK) is −0.0174 ($p = 0.228$). The denominator-free check therefore does not reveal a statistically significant fertilizer association. Because the FAOSTAT numerator represents national agricultural use across agricultural activities rather than coffee-specific fertilizer application [18], this exercise remains an aggregate measurement sensitivity analysis and is not interpreted as a coffee-farm dose–response estimate.

The auxiliary four-year harvested-area equation uses 336 country-year observations and country effects. Harvested area is highly persistent (rho = 0.9418, $p < 0.001$). The lagged international price coefficient is −0.0606 ($p = 0.005$). The negative sign is reported without imposing the theoretically expected direction; it shows that aggregate area adjustment in this sample is not a simple positive function of the common price signal. Most importantly, area is analyzed separately and no longer enters the yield equation.

Figure 3 shows that the sign and magnitude of individual long-run coefficients can change with sample composition. This sensitivity is consistent with the formal rejection of slope homogeneity and supports the decision not to privilege a single regional elasticity.

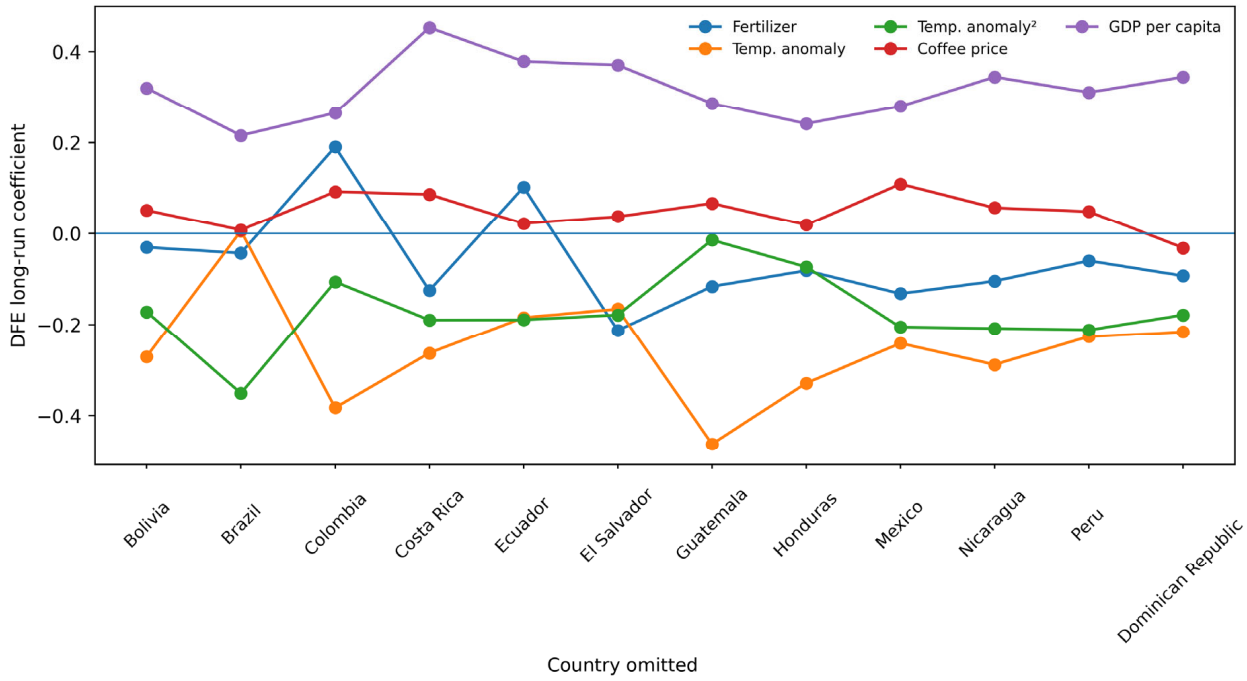

Figure 3. Leave-one-country-out DFE robustness for the no-area yield specification. Each line shows a long-run coefficient after omitting one country; DFE is used only as a sensitivity benchmark. In the legend, "Temp. anomaly" denotes the linear annual temperature-anomaly term, whereas "Temp. anomaly^2" denotes its squared term, included to capture nonlinear temperature effects.

Figure 4 shows that the residual CD statistics are not statistically significant for PMG, MG, or DFE. This supports the interpretation that residual common dependence is not statistically detectable after the reported dynamic specifications are estimated. Additional implementation and reproducibility details are provided in Appendix A.

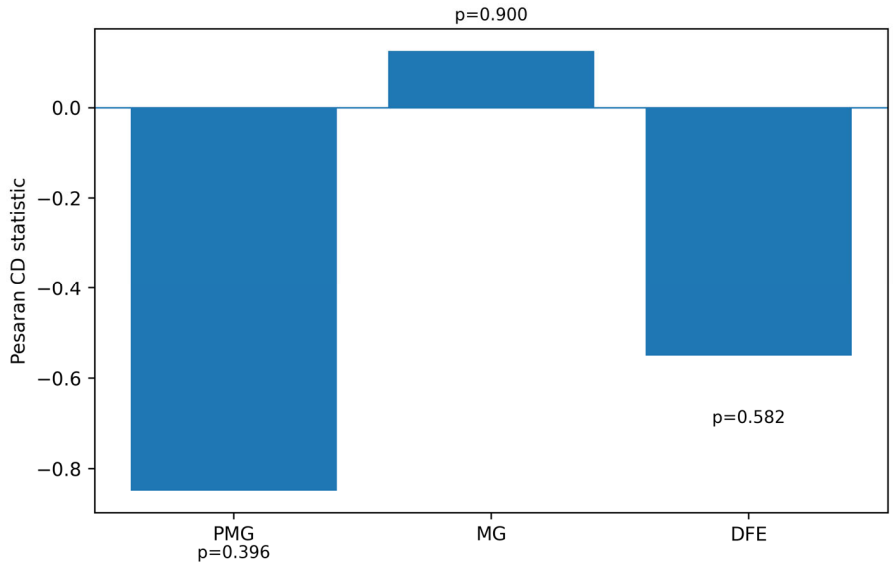

Figure 4. Residual Pesaran cross-sectional-dependence diagnostics for PMG, MG, and DFE for the no-area yield specification. All three *p*-values exceed 0.05.

4.6. Integrated Interpretation of the Results

Taken together, the evidence supports two robust conclusions: the ARDL(1,1) specifications contain negative and statistically significant dynamic-adjustment coefficients, and long-run slopes are heterogeneous across countries. The first result is interpreted

as conditional mean reversion rather than as proof of cointegration. By contrast, none of the constructed NPK-use/coffee-area proxy, temperature, price, or income coefficients is statistically significant at the 5% level in either central long-run specification, and no short-run mean coefficient is significant at 5%. The main contribution is therefore a disciplined characterization of dynamic adjustment and heterogeneity rather than a claim of a universal regional production elasticity.

The results are framed around H1 and H2. H1 is supported by the negative and significant adjustment coefficients in the ARDL(1,1) ECM parameterization, interpreted as dynamic mean reversion. The Westerlund statistics are not used to support H1 because the mixed I(0)/I(1) integration profile limits a cointegration interpretation in this setting. H2 is supported by the Swamy-type rejection of common long-run slopes. The evidence does not support a common statistically precise temperature or input coefficient at the national annual scale.

5. Discussion

The first substantive result is the dynamic adjustment itself. Both PMG and MG produce negative and highly significant mean adjustment coefficients in the ECM parameterization. Given that coffee yield, the constructed NPK-use/coffee-area proxy, and temperature are I(0), the result is interpreted as conditional mean reversion or dynamic adjustment, not as proof of a persistent cointegrating relationship. The standard Westerlund statistics are retained only as supplementary diagnostics—Gt, Pt, and Pa reject their nominal null, whereas Ga does not—reporting that disagreement, together with the mixed integration profile, is more informative than treating the procedure as definitive proof of cointegration.

The second result is structural heterogeneity rather than a harvested-area effect. After removing harvested area as a separate regressor from the yield equation, the central long-run coefficients on the constructed NPK-use/coffee-area proxy, temperature anomaly, international coffee price, and GDP per capita remain statistically imprecise. The Swamy-type test strongly rejects common slopes, indicating that country-specific regimes matter for interpretation. This finding is consistent with the literature emphasizing the sensitivity of coffee productivity to local agroecological and management conditions [4,11,13–16].

The fertilizer result is not robustly different from zero in either central estimator. The fertilizer metadata have now been independently verified against FAOSTAT: the numerator is national Agricultural Use of the product category 'NPK fertilizers' (HS 310520), measured in tonnes of product, and its coverage is national agriculture rather than coffee specifically [18]. Importantly, the targeted denominator-free DFE sensitivity check also yields an imprecise fertilizer association when national NPK agricultural use replaces the area-normalized proxy: the implied long-run coefficient is -0.0192 (SE = 0.0143, $p = 0.207$), and the short-run coefficient on the first difference of ln(1 + NPK) is -0.0174 ($p = 0.228$). Thus, the null fertilizer finding is not generated solely by the shared coffee-area denominator. Nevertheless, neither the area-normalized proxy nor the denominator-free numerator is a coffee-specific application measure. More detailed farm- or region-level nutrient data are required to identify nutrient composition, timing, soil status, application efficiency, or interactions with water and shade. Chéron-Bessou et al. [17] likewise show that fertilizer use can be environmentally important without implying a simple monotonic yield response.

The temperature terms are also statistically imprecise at the national annual scale. This does not overturn the substantial agroecological literature documenting climate sensitivity in coffee [12–14,16]. Instead, it highlights a scale mismatch: annual national anomalies cannot represent altitude, local extremes, shade, humidity, cultivar composition, or pheno-

logical timing. The study therefore does not infer a universal national temperature–response curve from coefficients whose confidence intervals include zero.

The international price and GDP per capita function as controls rather than central findings. The common international price is I(1) under the dedicated univariate tests and is included only in specifications without unrestricted year dummies. GDP per capita is not interpreted as a direct technological measure; it is a broad structural-development proxy. Neither coefficient is statistically significant in both central estimators.

A central methodological result is that the Swamy-type Wald test rejects long-run slope homogeneity. This is conceptually distinct from a Hausman test, and the distinction is maintained throughout the analysis [19]. MG therefore receives greater interpretive weight, but the cross-country mean remains a summary of heterogeneous national coefficients rather than a structural regional parameter.

The study has five explicit limitations. First, annual national data cannot represent subnational altitude, shade, soils, varieties, plantation age, irrigation, farm size, or management practices. Second, the national temperature anomaly is not specific to coffee-growing zones. Third, the constructed NPK-use/coffee-area proxy and yield may be jointly determined, so the dynamic model does not establish causal effects. Fourth, although the NPK numerator's unit and coverage are now documented from FAOSTAT, it measures national agricultural use of NPK fertilizer product in tonnes and is not coffee-specific [18]. Dividing this national numerator by coffee harvested area creates a scale-normalized proxy rather than an observed coffee-farm application rate; the denominator-free sensitivity check addresses the mechanical denominator issue but not this aggregation mismatch. Fifth, the common international price cannot be estimated together with unrestricted year fixed effects; the analysis addresses this identification constraint explicitly and uses a price-excluded year-fixed-effects robustness specification. These limitations define the scale of valid inference rather than negating the contribution of the dynamic regional comparison.

Future research should prioritize georeferenced or subnational panels, seasonal climate indicators, rainfall and drought extremes, plantation age, shade, soil quality, varieties, farm-level input management, and sources of exogenous variation. Separate species-specific models for *C. arabica* and *C. canephora* would be particularly valuable given their different climatic and agronomic profiles [7,8].

Policy implications should be proportionate to the evidence. The results support country-sensitive monitoring and adaptation rather than a uniform regional input prescription. Because the central slope coefficients are imprecise, the study does not infer that increasing fertilizer, area, or any single market incentive will raise yield by a common amount across Latin America.

At the program level, improving the quality of agronomic information may be more defensible than prescribing a uniform input quantity. Soil testing, nutrient management, climate monitoring, and locally adapted technical assistance are consistent with the broader literature, but they are not claimed as causal effects of the estimated national coefficients.

Climate adaptation remains important because the broader coffee literature documents substantial exposure to warming and shifting suitability [12–14]. The present macro-panel adds a different message: national annual temperature measures are too aggregated to recover those local biological mechanisms precisely, which strengthens the case for finer spatial data.

Finally, the separate area-adjustment exercise shows that planting and harvested-area dynamics are highly persistent and do not display a simple positive response to the common price series in this sample. This reinforces the decision to separate area adjustment from the yield function and to avoid attributing productivity changes mechanically to land expansion.

Overall, the evidence indicates statistically significant dynamic adjustment in the ARDL-ECM parameterization, heterogeneous conditional long-run slopes, and limited precision for common national-level climate, input, price, and income associations. The adjustment coefficients are interpreted as mean reversion rather than independent cointegration evidence, and the denominator-free NPK sensitivity check does not produce a statistically significant fertilizer association. This conclusion is aligned with the variable construction, the identification constraints, and the reported diagnostics.

6. Conclusions

This study provides a dynamic assessment of coffee yield in twelve Latin American countries over 1992–2023. Its main novelty is methodological and empirical: it separates yield from harvested-area adjustment, respects the collinearity between a common international price and unrestricted year effects, tests the price series as a genuine univariate time series, explicitly qualifies the Westerlund diagnostics under the mixed I(0)/I(1) integration profile, distinguishes PMG from standard MG correctly, reports both short- and long-run coefficients, documents the FAOSTAT NPK numerator's unit and national agricultural coverage, and adds a denominator-free NPK sensitivity check. The strongest findings are statistically significant dynamic adjustment or mean reversion in the ARDL-ECM parameterization and rejection of common long-run slopes; neither result is presented as causal evidence, and the adjustment coefficient is not treated as independent proof of cointegration. The data do not support a statistically precise common regional coefficient for the constructed NPK-use/coffee-area proxy, the denominator-free national NPK-use numerator, temperature anomaly, coffee price, or GDP per capita. Future research should therefore move toward subnational and georeferenced designs, species-specific panels, coffee-specific fertilizer-use measures, climate extremes, and identification strategies capable of supporting causal inference. In policy terms, the evidence favors country-sensitive monitoring and adaptation rather than uniform regional elasticities or input prescriptions.

Author Contributions: Conceptualization, L.R., R.-K.T.-M. and S.O.-M.; methodology, S.O.-M.; software, S.O.-M.; validation, L.R. and R.-K.T.-M.; formal analysis, L.R., R.-K.T.-M. and S.O.-M.; investigation, L.R., R.-K.T.-M. and S.O.-M.; resources, L.R. and R.-K.T.-M.; data curation, L.R. and R.-K.T.-M.; writing—original draft preparation, L.R., R.-K.T.-M. and S.O.-M.; writing—review and editing, L.R., R.-K.T.-M. and S.O.-M.; visualization, S.O.-M.; supervision, S.O.-M.; project administration, S.O.-M. S.O.-M. conceived the study, designed and implemented the econometric modeling strategy, directed the research process, conducted the formal analysis, interpreted the findings, and led the preparation of the original manuscript. L.R. and R.-K.T.-M. contributed to data collection and curation, the literature review, validation of the results, and critical revision of the manuscript. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The country-year dataset covers 12 Latin American coffee-producing countries over 1992–2023. The underlying data are publicly available from the Food and Agriculture Organization Corporate Statistical Database (FAOSTAT), the World Bank Climate Change Knowledge Portal (CCKP), the Climatic Research Unit (CRU), the World Development Indicators (WDI), and the International Monetary Fund (IMF) commodity-price series. The harmonized dataset and replication code are available in: https://doi.org/10.6084/m9.figshare.33368941.

Acknowledgments: The authors gratefully acknowledge the Department of Economics and the Vice-Rectorate for Research for their institutional support and continued commitment to promoting and strengthening research conducted by faculty members.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ARDL	Autoregressive Distributed Lag
ADF	Augmented Dickey-Fuller
CADF	Cross-sectionally Augmented Dickey-Fuller
CCKP	Climate Change Knowledge Portal
CIPS	Cross-sectionally Augmented IPS
CRU	Climatic Research Unit
DFE	Dynamic Fixed Effects
ECM	Error-Correction Model
FAOSTAT	Food and Agriculture Organization Statistical Database
GDP	Gross Domestic Product
ICO	International Coffee Organization
IMF	International Monetary Fund
KPSS	Kwiatkowski-Phillips-Schmidt-Shin
LR	Long run
MG	Mean Group
PMG	Pooled Mean Group
SD	Standard deviation
SR	Short run
WDI	World Development Indicators

Appendix A. Reproducibility and Robustness Details

The appendix summarizes implementation details that are directly relevant to reproducibility. It contains no additional substantive claims beyond the results reported in the main text.

Table A1. Reproducibility and robustness details for the empirical framework.

Check	Specification	Key Numerical Result	Interpretation
Lag selection	Common comparison sample; p,q <= 3	BIC ARDL(1,1) = 928.93; next-best ARDL(2,1) = 932.59	BIC selects ARDL(1,1)
Year-effects robustness	DFE; common price omitted; unrestricted year effects	ECM = −0.9305 ($p < 0.001$)	Adjustment remains significant without price/time-FE collinearity
Year-effects LR slopes	Same robustness specification	Fertilizer −0.1439 ($p = 0.666$); temp −0.4568 ($p = 0.429$); temp2 0.1588 ($p = 0.727$); GDPpc 0.5445 ($p = 0.349$)	No robust common LR slope
Area adjustment	Country effects; 4-year lags; N = 336	rho = 0.9418 ($p < 0.001$); price t − 4 = −0.0606 ($p = 0.0049$)	Area analyzed separately from yield
Residual CD	No-area yield models	PMG −0.848 ($p = 0.396$); MG 0.125 ($p = 0.900$); DFE −0.550 ($p = 0.582$)	Residual CD not statistically detected
Denominator-free fertilizer sensitivity	DFE-style ARDL(1,1); ln(1 + national NPK agricultural use in tonnes of product) replaces area-normalized proxy; country and unrestricted year effects; country-clustered inference [18]	N = 372; adjustment = −0.9266 ($p < 0.001$); LR NPK = −0.0192 (SE = 0.0143, $p = 0.207$); SR first-difference NPK = −0.0174 ($p = 0.228$)	No statistically significant fertilizer association without the shared harvested-area denominator; numerator remains national rather than coffee-specific

Note. ARDL = Autoregressive Distributed Lag; BIC = Bayesian Information Criterion; DFE = Dynamic Fixed Effects; ECM = error-correction-model parameterization; LR = long run; CD = Pesaran cross-sectional-dependence statistic. The common price is omitted whenever unrestricted year fixed effects are included. The denominator-free NPK sensitivity check uses country-clustered inference; the reported long-run NPK standard error is based on the delta method. The NPK numerator is FAOSTAT Fertilizers by Product, item 'NPK fertilizers' (HS 310520), element Agricultural Use, measured in tonnes of product; its scope is national agricultural use and is not coffee-specific [18].

References

1. Poncet, V.; van Asten, P.; Millet, C.P.; Vaast, P.; Allinne, C. Which diversification trajectories make coffee farming more sustainable? *Curr. Opin. Environ. Sustain.* **2024**, *68*, 101432. [CrossRef]
2. Baffes, J.; Lewin, B.; Varangis, P. Coffee: Market setting and policies. In *Commodity Market Reforms in Africa*; Varangis, P., Lewin, B., Eds.; World Bank: Washington, DC, USA, 2005; pp. 297–309.
3. Giovannucci, D.; Potts, J.; Killian, B.; Wunderlich, C.; Schuller, S.; Henschke, C.; Pinard, F. *Sustainable Coffee: A Guide to the World of Coffee and Sustainability*; International Institute for Sustainable Development: Winnipeg, MB, Canada, 2020.
4. Bilen, C.; El Chami, D.; Mereu, V.; Trabucco, A.; Marras, S.; Spano, D. A systematic review on the impacts of climate change on coffee agrosystems. *Plants* **2023**, *12*, 102. [CrossRef]
5. Nicholson, W.; Snyder, C. *Microeconomic Theory: Basic Principles and Extensions*, 11th ed.; Cengage Learning: Boston, MA, USA, 2015.
6. Pesaran, M.H.; Shin, Y.; Smith, R.P. Pooled mean group estimation of dynamic heterogeneous panels. *J. Am. Stat. Assoc.* **1999**, *94*, 621–634. [CrossRef]
7. Jan-Smith, E.; Downes, H.; Davis, A.P.; Richard-Bollans, A.; Haggar, J.; Sarmu, D.; Kite, G.C.; Howes, M.-J.R. Metabolomic insights into the Arabica-like flavour of stenophylla coffee and the chemistry of quality coffee. *npj Sci. Food* **2025**, *9*, 33. [CrossRef]
8. World Coffee Research. *The Roots of Robusta: Cultivating Growth for a Species Once Overlooked*; World Coffee Research: Portland, OR, USA, 2024.
9. Davis, A.P.; Shepherd-Clowes, A.; Cheek, M.; Moat, J.; Luo, D.W.; Kiwuka, C.; Kalema, J.; Tchiengue, B.; Viruel, J. Genomic data define species delimitation in Liberica coffee with implications for crop development and conservation. *Nat. Plants* **2025**, *11*, 1729–1738. [CrossRef]
10. Kouadio, L.; Byrareddy, V.M.; Sawadogo, A.; Newlands, N.K. Probabilistic yield forecasting of Robusta coffee at the farm scale using agroclimatic and remote sensing-derived indices. *Agric. For. Meteorol.* **2021**, *306*, 108449. [CrossRef]
11. Della Peruta, R.; Mereu, V.; Spano, D.; Marras, S.; Vezy, R.; Trabucco, A. Projecting trends of Arabica coffee yield under climate change: A process-based modelling study at continental scale. *Agric. Syst.* **2025**, *227*, 104353. [CrossRef]
12. Bunn, C.; Läderach, P.; Rivera, O.O.; Kirschke, D. A bitter cup: Climate change profile of global production of Arabica and Robusta coffee. *Clim. Change* **2015**, *129*, 89–101. [CrossRef]
13. Ovalle-Rivera, O.; Läderach, P.; Bunn, C.; Obersteiner, M.; Schroth, G. Projected shifts in *Coffea arabica* suitability among major global producing regions due to climate change. *PLoS ONE* **2015**, *10*, e0124155. [CrossRef]
14. Grüter, R.; Trachsel, T.; Laube, P.; Jaisli, I. Expected global suitability of coffee, cashew and avocado due to climate change. *PLoS ONE* **2022**, *17*, e0261976. [CrossRef]
15. Byrareddy, V.M.; Kath, J.; Kouadio, L.; Mushtaq, S.; Geethalakshmi, V. Assessing scale-dependency of climate risks in coffee-based agroforestry systems. *Sci. Rep.* **2024**, *14*, 8028. [CrossRef]
16. Craparo, A.C.W.; van Asten, P.J.A.; Läderach, P.; Jassogne, L.T.P.; Grab, S.W. *Coffea arabica* yields decline in Tanzania due to climate change: Global implications. *Agric. For. Meteorol.* **2015**, *207*, 1–10. [CrossRef]
17. Chéron-Bessou, C.; Acosta-Alba, I.; Boissy, J.; Payen, S.; Rigal, C.; Setiawan, A.A.R.; Sevenster, M.; Tran, T.; Azapagic, A. Unravelling life cycle impacts of coffee: Why do results differ so much among studies? *Sustain. Prod. Consum.* **2024**, *47*, 251–266. [CrossRef]
18. Food and Agriculture Organization of the United Nations (FAO). *FAOSTAT Domain Fertilizers by Product: Methodological Note, Release July 2025*; FAO: Rome, Italy, 2025. Available online: https://files-faostat.fao.org/production/RFB/RFB_EN_README.pdf (accessed on 3 September 2026).
19. Swamy, P.A.V.B. Efficient inference in a random coefficient regression model. *Econometrica* **1970**, *38*, 311–323. [CrossRef]
20. Pesaran, M.H. A simple panel unit root test in the presence of cross-section dependence. *J. Appl. Econom.* **2007**, *22*, 265–312. [CrossRef]
21. Dickey, D.A.; Fuller, W.A. Distribution of the estimators for autoregressive time series with a unit root. *J. Am. Stat. Assoc.* **1979**, *74*, 427–431. [CrossRef]
22. Kwiatkowski, D.; Phillips, P.C.B.; Schmidt, P.; Shin, Y. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *J. Econom.* **1992**, *54*, 159–178. [CrossRef]
23. Pesaran, M.H. *General Diagnostic Tests for Cross-Section Dependence in Panels*; (Cambridge Working Papers in Economics No. 0435); University of Cambridge: Cambridge, UK, 2004.
24. Westerlund, J. Testing for error correction in panel data. *Oxf. Bull. Econ. Stat.* **2007**, *69*, 709–748. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.