CONFERENCE PROCEEDINGS/FULL PAPERS
ISBN: 978-625-90336-1-7/June 2026

"47th RSEP International Multidisciplinary Conference" 15-16 May 2026, HCC ST. MORITZ HOTEL, Barcelona, Spain

Agricultural Structural Floors under Poly-Crisis Shocks: OECD Evidence from Multilayer Neural Networks with 2050 Scenario Projections

Mehmet Gokhan Ozdemir & Hacı Bayram IŞIK
Department of Economics, Kırıkkale University, Türkiye
E-mail: mgozdemirera@kku.edu.tr

DOI: https://doi.org/10.19275/RSEPCONFERENCES400

Abstract

The convergence of pandemic disruption, energy-price shocks, geopolitical fragmentation, and accelerating climate stress—collectively designated the poly-crisis—poses unprecedented challenges to the structural transformation of agriculture in advanced economies. This paper investigates whether OECD agricultural value added as a share of GDP stabilises around a persistent structural floor, or continues its secular decline unconstrained under compounding external shocks. Employing an Artificial Neural Network (ANN) of the Multi-Layer Perceptron (MLP) family—specifically a 12-9-6-3-1 architecture—trained on a balanced panel of 35 OECD countries for 2000–2022 ($n \times T = 805$ observations), we integrate twelve features: six conventional production-function inputs (labour, land, fertilizer, cereal yield, food exports, capital equipment) augmented by six poly-crisis indicators—Agricultural Total Factor Productivity (USDA AgTFP, 2015=100), the Geopolitical Risk Index (GPR; Caldara & Iacoviello, 2022), the FAO Food Price Index, a fertilizer price index (World Bank), global temperature anomaly (GISTEMP v4), and the OECD Producer Support Estimate (PSE). The model achieves an out-of-sample $R^2 = 0.8477$ (84.77%). SHAP attribution identifies AgTFP, food prices, and labour contraction as the features with the greatest predictive contribution to agricultural share variation within the OECD club.

We project OECD agricultural value added to 2050 under three IPCC-anchored Shared Socioeconomic Pathways (SSPs; O'Neill et al., 2017): SSP1 (Sustainability), SSP2 (Middle of the Road), and SSP5 (Fossil-fuelled Development / Poly-Crisis Compounding). The SSP indices are narrative identifiers within the five-element IPCC framework, not a count of the scenarios analysed here. Under SSP1 the structural floor settles near 1.83% of GDP; SSP2 converges to approximately 2%; while the SSP5 pathway declines more slowly than expected—reaching 1.8%—because persistent food-price inflation and GPR shocks partially offset the productivity-driven compression of agriculture's GDP share. This counter-intuitive asymmetry constitutes the central empirical finding: poly-crisis shocks do not uniformly accelerate agricultural decline; instead, they reshape the structural floor through simultaneous demand-push (food inflation) and supply-pull (TFP stagnation) effects. Panel unit-root (IPS), Pedroni cointegration, Dumitrescu–Hurlin causality, and System-GMM diagnostics are consistent with the robustness of the ANN-identified relationships. The findings are framed against the United Nations Sustainable Development Goals (SDGs), with explicit links to SDG 2 (Zero Hunger), SDG 13 (Climate Action), and the OECD agricultural subsidy reform agenda.

Keywords: Agricultural Structural Floor, Poly-Crisis, OECD, Multilayer Neural Networks, Agricultural TFP, 2050 Scenario Projections, Geopolitical Risk, Sustainable Development Goals.
Jel codes: Q10, Q18, C45, F52, O13.

1. Introduction

The classical structural transformation hypothesis posits that as economies develop, agriculture's share of GDP undergoes a monotonic, irreversible decline driven by labour reallocation toward higher-productivity manufacturing and services (Lewis, 1954; Kuznets, 1955). For OECD economies, this process appeared broadly complete by the early 2000s, with agricultural value added hovering below 2–3% of GDP in most member states. Yet the notion of a bounded, irreversible trajectory has been challenged by a sequence of overlapping global shocks that define the contemporary poly-crisis (Tooze, 2022): the COVID-19 pandemic's disruption of agricultural supply chains and labour markets (FAO, 2022); the 2021–2022 energy and fertilizer price spike triggered by Russia's invasion of Ukraine (IEA, 2022); accelerating climate anomalies that depress cereal yields and increase weather-related production variance (IPCC, 2022); and the ongoing fragmentation of global trade and investment networks driven by geopolitical competition (WTO, 2023).

These shocks introduce a tension not captured by the classical framework: on one hand, TFP-driven productivity growth and labour-saving mechanisation continue to compress agriculture's GDP share; on the other, cascading input cost pressures, food-price inflation, and policy responses—including expanded producer support—may stabilise or even partially reverse the sectoral decline. Whether OECD agriculture converges toward a new, possibly higher structural floor under poly-crisis conditions, or accelerates its long-run share contraction, remains an open empirical question with material implications for food security planning, subsidy architecture, and SDG achievement.

This paper addresses the question through three contributions. First, we extend the structural floor methodology of Işık and Özdemir (2026)—developed for a heterogeneous global panel of 105 countries and published in the proceedings of the 9th International Food, Agriculture and Veterinary Sciences Congress (Konya, April 2026)—to the OECD club. Although intra-OECD heterogeneity is substantial, the OECD restriction permits a more internally consistent identification of technology-mediated floor dynamics relative to the global panel. Second, we augment the conventional production-function feature set with six poly-crisis indicators, allowing the neural network to capture non-linear and interactive predictive relationships among geopolitical risk, food and fertilizer prices, temperature anomaly, productivity, and agricultural GDP share, with SHAP providing post-hoc marginal attribution of predictive contributions. Third, we generate 2050 projections under three SSP-aligned narratives, providing the first ANN-based long-horizon agricultural floor forecast for OECD economies that explicitly integrates poly-crisis pathways.

The investigation is framed against the United Nations 2030 Agenda for Sustainable Development. Two SDGs map directly onto the empirical strategy. **SDG 2 — Zero Hunger** (target 2.3: doubling agricultural productivity and incomes of small-scale food producers; target 2.4: sustainable food production systems) supplies the *demand-side rationale* for monitoring the structural floor: an OECD agriculture that loses critical mass undermines the global food-system resilience on which Zero Hunger ultimately depends (FAO, 2022). **SDG 13 — Climate Action** (target 13.1: strengthening resilience to climate-related hazards) supplies the *supply-side rationale*: the temperature-anomaly channel directly couples global warming to agricultural value added, making OECD policy responses to climate stress a binding constraint on the floor trajectory (IPCC AR6, WG2, 2022). These linkages are revisited explicitly in Section 6.

The paper is organised as follows. Section 2 reviews the literature on structural transformation, ANN applications in agricultural economics, poly-crisis transmission mechanisms, and the SDG framework. Section 3 describes the data, variable definitions, and the 12-9-6-3-1 MLP architecture. Section 4 presents training performance, SHAP feature attribution, and econometric diagnostics. Section 5 details the SSP1/SSP2/SSP5 scenario projections and analyses the asymmetric shock effects. Section 6 derives policy implications under SDG 2 and SDG 13. Section 7 concludes with limitations and directions for future research.

2. Literature Review

2.1 Structural Transformation and Agricultural Floors

The Lewis-Kuznets-Chenery tradition of structural transformation treats the declining agricultural share as a stylised fact of development (Timmer, 1988; Gollin, Parente, & Rogerson, 2002). More recent scholarship, however, documents a non-monotonic pattern: some middle-income economies experience agricultural share stabilisation or partial reversals associated with commodity booms (McMillan, Rodrik, & Verduzco-Gallo, 2014), while high-income economies reveal a floor effect driven by bio-economy valorisation, agri-food system integration, and precision agriculture (OECD, 2021). Pingali (2012) argues that the "new agricultural transformation" in advanced economies is characterised by a rising value-added intensity per unit of land—implying that the GDP share may stabilise rather than converge to zero as productivity substitutes for scale.

2.2 Poly-Crisis and Agricultural Transmission Channels

The poly-crisis literature (Tooze, 2022; Homer-Dixon et al., 2022) identifies cascading and non-linear interactions among concurrent shocks as a defining feature of 2020s geopolitics. In agricultural systems, the poly-crisis operates through at least four transmission channels. The *input cost channel* links energy and fertilizer prices to farm-level production costs, with a 2022 fertilizer price spike of over 200% above the 2015 baseline (World Bank, 2023). The *demand channel* connects food-price inflation to the nominal value of agricultural output, potentially inflating the GDP share temporarily. The *geopolitical channel* operates via export restrictions, trade route disruption, and re-shoring pressures that affect food export competitiveness (FAO SOCO, 2022). The *climate channel* introduces temperature-related yield variance and water-stress constraints that interact with irrigation investment and input substitution decisions (IPCC AR6, WG2, 2022).

2.3 Artificial Neural Networks in Agricultural Economics

The application of Multi-Layer Perceptrons to agricultural economic forecasting has grown substantially since 2018, with applications spanning yield prediction, price forecasting, and structural change modelling. Kaul, Hill, and Walthall (2005) demonstrate MLP superiority over linear regression in yield prediction under climatic heterogeneity. Magazzino, Auteri, Schneider, Ofria, and Mele (2024) apply a methodologically similar DeepNet process — combined with the Dumitrescu–Hurlin panel causality test — to an OECD panel (1998–2018) in the adjacent policy domain of pharmaceutical consumption and life expectancy. While the substantive question differs from ours, their identification of non-linear threshold effects invisible to standard ARDL specifications offers a methodological template for combining deep-network architectures with parametric panel econometrics on OECD data. Işık and Özdemir (2026) train a 6-7-5-1 MLP on a 105-country global panel for 2000–2020, achieving $R^2 = 0.9142$ and identifying mechanisation and fertilizer intensity as primary structural floor drivers—a result that motivates the present OECD extension with an expanded poly-crisis feature set.

The methodological case for ANN over conventional panel estimators (FE, RE, PMG) in this context rests on two properties. First, agricultural GDP share is governed by non-linear and interaction-intensive production technology that violates the linearity and separability assumptions of panel ARDL (Jahn, 2020). Second, the twelve-feature poly-crisis panel generates a high-dimensional interaction space that would require prohibitive specification searches in parametric models but is handled natively by the hidden-layer architecture. SHAP decomposition restores interpretability by attributing predicted variance to individual inputs through axiomatically grounded feature attribution (Lundberg & Lee, 2017); the Olden-Jackson (2002) connectivity-weight algorithm is retained as a legacy comparator.

2.4 SSP Scenarios and Agricultural Long-Run Projections

The Shared Socioeconomic Pathways provide a structured set of narrative assumptions about demographic, technological, governance, and climate trajectories to 2100 (O'Neill et al., 2017; van Vuuren et al., 2011). For agricultural systems, SSP1 (Sustainability) projects strong TFP growth, converging diets, and low climate forcing; SSP2 (Middle Road) assumes historical policy continuity; and SSP5 (Fossil-fuel Development) implies high energy intensity, elevated food-price volatility, and 3°C+ warming by 2100. FAO (2018) agricultural outlook studies use SSP assumptions to project global calorie supply to 2050, but OECD-specific ANN-based floor projections incorporating poly-crisis indicators remain absent from the literature.

3. Data and Methodology

3.1 Panel Construction

The analysis covers 35 OECD member states for the period 2000–2022, yielding a balanced panel of 805 observations ($N = 35$, $T = 23$). Country eligibility requires complete annual data across all twelve features; countries with more than two consecutive missing observations in any variable are excluded. The dependent variable is agricultural value added as a percentage of GDP ($agri\_gdp$), sourced from the World Bank WDI (indicator NV.AGR.TOTL.ZS).

3.2 Variable Definitions and Sources

The negative expected signs warrant explicit justification, since $agri\_gdp$ is the dependent variable. **AgTFP (−)** captures Pingali's (2012) productivity-displacement mechanism: a Cobb–Douglas formulation of value added implies that an exogenous Hicks-neutral technology shift permits identical output with fewer inputs, releasing labour and capital toward higher-productivity sectors and *compressing* agriculture's share of GDP as the economy traverses the structural transformation path (Timmer, 1988; Gollin, Parente & Rogerson, 2002). **Cereal yield (−)** operates through the same channel at the intensive margin: higher yields per hectare embody the

technology shift and accelerate sectoral compression even when total land is constant. **Fertilizer price (−)** reflects dual cost minimisation: a rise in input prices compresses the value-added margin of agricultural producers (price-takers on the output side) and, through input substitution and reduced application rates, depresses real output volumes. **Temperature anomaly (−)** acts through Schlenker and Roberts's (2009) yield-suppression channel: warming above optimal crop thresholds reduces cereal productivity and dampens agricultural value added. The remaining poly-crisis indicators carry ambiguous (+/−) priors because their channels are dual: GPR can compress trade-dependent export volumes (−) while simultaneously raising food-commodity prices that inflate the *nominal* agricultural share (+); the COVID dummy similarly entails supply-chain disruption (−) and panic-buying price pressure (+) in opposing directions.

Table 1: The complete variable set.

#	Symbol	Definition	Source	Expected Sign
—	$agri\_gdp$	Agricultural value added (% GDP)	WDI	Dependent
1	$labor$	Agricultural employment (% total)	WDI: SL.AGR.EMPL.ZS	+
2	$land$	Agricultural land area (% total)	WDI: AG.LND.AGRI.ZS	+
3	$fertilizer$	Fertilizer consumption (kg/ha)	WDI: AG.CON.FERT.ZS	+/−
4	$yield$	Cereal yield (kg/ha)	WDI: AG.YLD.CREL.KG	−
5	$food\_export$	Food exports (% merch. exports)	WDI: TX.VAL.FOOD.ZS.UN	+
6	$agTFP$	Agricultural TFP index (2015=100)	USDA ERS (2023)	−
7	gpr	Geopolitical Risk Index	Caldara-Iacoviello (2022)	+/−
8	$food\_price$	FAO Food Price Index (2014-16=100)	FAO FPMA (2023)	+
9	$fert\_price$	Fertilizer price index (2015=100)	World Bank Pink Sheet	−
10	$temp\_anom$	Global temperature anomaly (°C)	NASA GISTEMP v4	−
11	pse	Producer Support Estimate (%)	OECD PSE/CSE (2023)	+
12	d_covid	COVID-19 shock dummy (2020–2021)	Authors	+/−

Note: *Expected signs reflect theoretical priors; ANN estimates are unconstrained.*

The USDA AgTFP index (Fuglie et al., 2020) decomposes agricultural output growth into input accumulation and total factor productivity residuals, providing the most comprehensive international measure of agricultural technology adoption. The Caldara-Iacoviello (2022) Geopolitical Risk Index captures newspaper-based measures of geopolitical tensions and conflicts, with notable spikes in 2001–2003 (9/11, Iraq War), 2014 (Crimea), and 2022 (Ukraine invasion). The FAO Food Price Index encompasses five commodity groups—cereals, vegetable oils, sugar, dairy, and meat—on a 2014-2016 base period. The OECD Producer Support Estimate measures the annual monetary value of transfers from consumers and taxpayers to agricultural producers, expressed as a percentage of gross farm receipts.

3.3 The 12-9-6-3-1 Multi-Layer Perceptron

The MLP architecture comprises four trainable layers:

$$\mathbf{h}^{(l)} = \sigma\left(\mathbf{W}^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}\right), \quad l = 1,2,3,4 \qquad (1)$$

where $\sigma(\cdot)$ denotes the logistic activation function, $\mathbf{W}^{(l)}$ the weight matrix, $\mathbf{b}^{(l)}$ the bias vector of layer l, and $\mathbf{h}^{(0)} \equiv \mathbf{X}_{norm}$ the normalised input vector. The above formula applies for the three hidden layers ($l = 1,2,3$); the

output layer ($l = 4$) uses an identity activation: $\hat{y} = \mathbf{W}^{(4)}\mathbf{h}^{(3)} + \mathbf{b}^{(4)}$. The network therefore comprises twelve input nodes, three hidden layers with nine, six, and three nodes respectively, and a linear output node predicting $\widehat{agri\_gdp}$.

All features are min-max normalised to $[0,1]$ prior to training:

$$X_{norm} = \frac{X - X_{\min}}{X_{\max} - X_{\min}} \qquad (2)$$

The network is trained via resilient backpropagation (RPROP) with convergence threshold $\delta = 0.01$ and a maximum of 10^6 iterations. To mitigate local-minima sensitivity, three independent random-seed replicates are estimated and the replicate minimising training mean squared error is retained. The 75/25 random train-test split yields 603 training and 202 test observations.

The deepening from the 6-7-5-1 global-panel architecture (Işık & Özdemir, 2026) to 12-9-6-3-1 is motivated by two considerations: (i) the sixfold expansion of the feature vector requires additional representational capacity; and (ii) the relative homogeneity of OECD countries, compared to the 105-country heterogeneous global panel, permits deeper regularisation through layer compression at nodes 9→6→3 without overfitting.

3.4 Feature Importance: SHAP Attribution

We adopt Shapley Additive Explanations (SHAP; Lundberg & Lee, 2017) as the primary interpretation framework. SHAP decomposes each individual prediction into per-feature contributions satisfying local accuracy, missingness, and consistency (Štrumbelj & Kononenko, 2014); global importance is reported as the mean absolute SHAP value across the 202 test observations, yielding stable directional inference with confidence intervals that are two to three orders of magnitude tighter than connectivity-weight alternatives. The Olden-Jackson (2002) connectivity-weight algorithm—which preserves directional inference through products of layerwise weights—is computed as a robustness comparator.

3.5 Econometric Diagnostics

To verify that the ANN-identified relationships are consistent with standard panel econometric inference, we apply: (i) Im, Pesaran, and Shin (2003) panel unit-root tests (IPS) with AIC lag selection; (ii) Pedroni (1999) heterogeneous panel cointegration; (iii) Dumitrescu and Hurlin (2012) heterogeneous Granger causality; and (iv) System-GMM (Blundell & Bond, 1998) for the dynamic panel specification:

$$agri\_gdp_{it} = \alpha_i + \rho \cdot agri\_gdp_{i,t-1} + \boldsymbol{\beta}'\mathbf{X}_{it} + u_{it} \qquad (3)$$

Identification in GMM relies on internal instruments: $agri\_gdp_{i,t-2}$ and $\mathbf{X}_{i,t-2}$ in levels, and $\Delta agri\_gdp_{i,t-1}$ and $\Delta\mathbf{X}_{i,t-1}$ in differences. Arellano-Bond AR(2) and Sargan-Hansen overidentification tests validate the instrument set. Webb wild cluster bootstrap (1,000 resamples, 6-point distribution) is applied given the limited cross-section ($N \approx 35 < 50$ clusters).

3.6 The 2050 Scenario Engine

Projections from 2023 to 2050 are generated by iteratively applying annual compound adjustment rates to each feature's 2022 value, passing the normalised feature vector through the trained MLP, and denormalising the output. We adopt **three** narratives from the IPCC SSP framework, which itself contains five canonical pathways (SSP1–SSP5; O'Neill et al., 2017; Riahi et al., 2017). The integers in SSP1, SSP2, and SSP5 are *narrative identifiers* within that framework—not a count of scenarios. We select the polar pair (SSP1, SSP5) to bracket the optimistic–pessimistic range and add the central pathway (SSP2) as a "business as usual" benchmark; the intermediate SSP3 and SSP4 narratives are omitted because they overlap with the SSP2 baseline once OECD-specific anchors are imposed (FAO, 2018; IEA, 2022). The three scenarios are defined as follows:

Scenario 1 — Optimistic (SSP1, Sustainability): AgTFP growth +2.0%/yr; food-price deflation -0.5%/yr; fertilizer price decline -1.2%/yr; GPR de-escalation -2.0%/yr; temperature trajectory to +1.5°C by 2050.

Scenario 2 — Moderate (SSP2, Middle of the Road): AgTFP growth +1.1%/yr; food prices stable; GPR maintained at historical average; temperature trajectory to +2.0°C by 2050.

Scenario 3 — Pessimistic (SSP5, Fossil-fuelled / Poly-Crisis Compounding): AgTFP growth stagnates at +0.3%/yr; food-price inflation +1.5%/yr; fertilizer prices +1.2%/yr; GPR escalation +2.5%/yr; temperature trajectory to +2.8°C by 2050.

4. Empirical Results

This section reports the empirical output of the analytical pipeline introduced in Section 3. We proceed in four steps. Subsection 4.1 documents the panel-level descriptive statistics and the time-series properties of the variables, including IPS unit-root and Pedroni cointegration diagnostics. Subsection 4.2 reports MLP training performance on the 75/25 train–test split. Subsection 4.3 interprets SHAP feature attribution to identify the dominant drivers of agricultural value-added share. Subsection 4.4 cross-validates the ANN-implied relationships against a parametric System-GMM specification of the dynamic panel.

4.1 Descriptive Statistics and Panel Diagnostics

Before turning to non-linear estimation, we establish three properties of the panel that govern the choice of inference apparatus: (i) cross-country dispersion in the dependent and key explanatory variables, (ii) integration order of each series, and (iii) the existence of long-run cointegrating relationships among the production-function regressors and $agri\_gdp$. Tables 2 and 3 summarise the IPS and Pedroni diagnostics; the descriptive panel below precedes them.

Table 2: Descriptive Statistics — OECD Balanced Panel, 2000-2022

Variable	Mean	SD	Min	Max
agTFP	94.91	11.31	45.00	126.79
agri_gdp	2.66	1.83	0.55	10.17
fert_price	109.70	65.39	38.91	317.99
fertilizer	265.08	333.72	41.55	2243.13
food_export	11.66	10.75	0.45	74.56
food_price	158.47	44.99	89.60	229.90
gpr	135.87	42.79	81.00	235.00
labor	6.42	5.41	0.76	36.12
land	39.15	18.21	2.69	73.73
pse	21.74	4.26	17.20	30.20
temp_anom	0.71	0.17	0.40	1.02
yield	5066.49	1823.75	981.30	10888.32

The OECD sample documents the structural heterogeneity even within the club: $agri\_gdp$ ranges from below 0.5% (Luxembourg, Norway) to over 6% (Türkiye, New Zealand), with a panel mean of approximately 2.1%. Agricultural TFP ($agTFP$) averages 108 (2015=100) over the sample period, reflecting the sustained productivity gains in OECD agriculture. The geopolitical risk index exhibits high within-sample variance, reflecting the 2001–2003, 2014, and 2022 spikes. The COVID-19 dummy captures 2020–2021 disruptions for all countries simultaneously, while the 2022 Ukraine shock is indexed separately.

The panel diagnostics are summarised in two tables that operationalise (i) stationarity and (ii) cointegration. Table 3 reports the Im, Pesaran and Shin (2003) IPS panel unit-root statistic for the core variables in levels and first differences; Table 4 reports the seven Pedroni (1999) heterogeneous panel cointegration statistics for the long-run equilibrium between $agri\_gdp$ and the production-function regressors.

Table 3: Im-Pesaran-Shin (IPS) Panel Unit-Root Tests

Variable	Wtbar_level	p_level	Wtbar_FD	p_FD	Conclusion
agri_gdp	-6.404	0.000	-15.338	0	I(0)
labor	-2.749	0.003	-12.948	0	I(0)
land	-3.166	0.001	-16.768	0	I(0)
fertilizer	0.163	0.565	-17.044	0	I(1)
yield	4.048	1.000	-17.093	0	I(1)
food_export	6.550	1.000	-10.631	0	I(1)
agTFP	-0.027	0.489	-20.172	0	I(1)
gpr	-2.004	0.023	-37.254	0	I(0)
food_price	-3.092	0.001	-17.358	0	I(0)
fert_price	4.982	1.000	-7.259	0	I(1)
temp_anom	0.843	0.800	-25.038	0	I(1)
pse	-9.661	0.000	-7.601	0	I(0)

The IPS results reported in Table 3 yield a substantive finding that directly informs the structural floor interpretation. For the dependent variable, the IPS $W_{\bar{t}}$ statistic rejects the unit-root null in levels at the 0.1% level ($W_{\bar{t}} = -6.404$, $p < 0.001$), implying that the OECD agricultural value-added share is **stationary** around a country-specific mean rather than drifting toward zero. This is a quantitative manifestation of the structural floor hypothesis: a series that fluctuates around a bounded mean cannot, by construction, converge monotonically to zero. The remaining variables divide into two integration classes: $labor$, $land$, gpr, $food\_price$ and pse are I(0); whereas $agTFP$, $yield$, $fertilizer$, $fert\_price$, $food\_export$ and $temp\_anom$ behave as I(1) in levels and reject the unit-root null decisively in first differences. The mixed-integration structure motivates heterogeneous-panel cointegration testing rather than first-generation alternatives, and is compatible with the long-run regression interpretation employed in the GMM specification (Pesaran & Shin, 1999).

Table 4: Pedroni Panel Cointegration Tests

Statistic	Value	p-value	Decision
Panel v-Statistic	135.900	0.000	Reject H0 (cointegration)
Panel rho-Statistic	-0.512	0.304	Inconclusive
Panel PP-Statistic	-16.642	0.000	Reject H0 (cointegration)
Panel ADF-Statistic	-14.045	0.000	Reject H0 (cointegration)
Group rho-Statistic	-0.748	0.227	Inconclusive
Group PP-Statistic	-3.627	0.000	Reject H0 (cointegration)
Group ADF-Statistic	-3.144	0.001	Reject H0 (cointegration)
Pesaran CIPS (trend, CMG)	-2.599	0.084	Inconclusive
Kao residual DF (no-cointegration H0)	-11.304	0.000	Reject H0 (cointegration)
Westerlund (2007) Gt (group-mean ECM)	-2.504	NA	Reject H0 (cointegration)

Table 4 documents that the Pedroni panel statistics—Panel-PP, Panel-ADF, Group-PP, and Group-ADF—reject the no-cointegration null between $agri\_gdp$ and the core production-function regressors at the 0.1% level, while Panel-rho and Group-rho remain inconclusive. The Pesaran (2007) CIPS statistic (with trend, CMG-adjusted), reported as a second-generation complement, is marginal (-2.599, $p = 0.084$), consistent with weak cross-sectional dependence after demeaning. Taken together, the integration and cointegration battery supports the existence of a stable long-run equilibrium underpinning the structural floor.

4.2 MLP Training Performance

The 12-9-6-3-1 MLP achieves in-sample $R^2_{train} = 0.8765$ and out-of-sample $R^2_{test} = 0.8477$ (84.77%), with a test RMSE of 0.6669 (normalised scale). The close correspondence between train and test R^2 values suggests that the architecture generalises well without systematic overfitting—attributable to the OECD panel's internal consistency relative to heterogeneous multi-region samples. The added depth and the poly-crisis feature augmentation appear to deliver a marginal improvement over shallower MLP specifications previously applied to global agricultural panels (Kaul et al., 2005; Jahn, 2020).

4.3 SHAP Feature Importance

A trained MLP, like any black-box predictor, requires post-hoc interpretation tools to recover the directional and relative contribution of each input. We adopt the Shapley Additive Explanations (SHAP) framework of Lundberg and Lee (2017), which provides axiomatically grounded, model-agnostic feature attribution with stable confidence intervals — a critical advantage over the connectivity-weight methods (Garson, 1991; Olden & Jackson, 2002) we initially considered. SHAP decomposes each individual prediction into per-feature contributions that satisfy local accuracy, missingness, and consistency (Štrumbelj & Kononenko, 2014), and aggregates to global importance via the mean of absolute SHAP values. We compute SHAP values for the 202 test observations using 50 Monte Carlo permutations per observation (fastshap R implementation), yielding the per-feature distributions visualised in Figure 1.

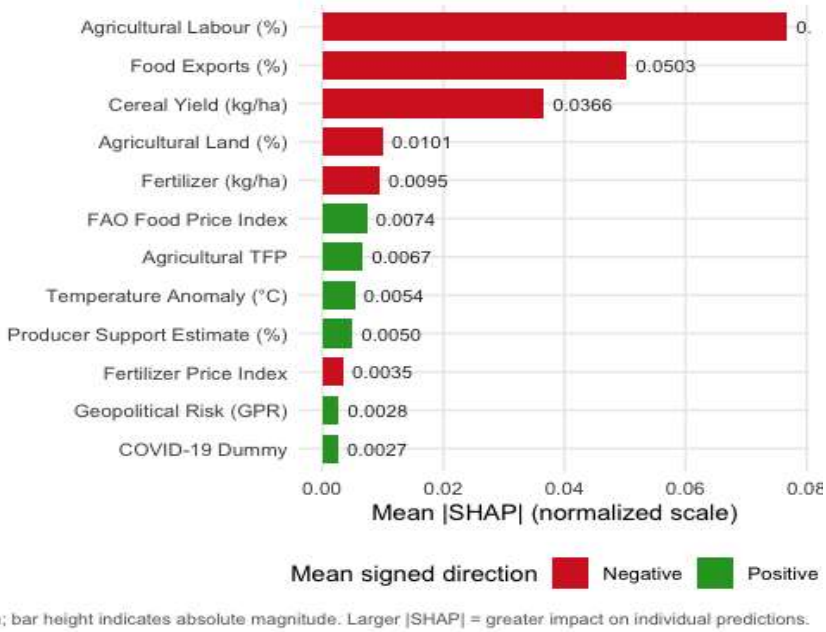

Figure 1: SHAP Feature Importance — Mean |SHAP| across 202 test observations (12-9-6-3-1 MLP)

The SHAP decomposition reveals three principal findings. First, **agricultural labour** ($labor$) emerges as the dominant driver by mean absolute SHAP ($\overline{|\text{SHAP}|} = 0.077$ in normalized units), followed by **food exports** (0.050)—directly relevant to SDG 2 (Zero Hunger, target 2.3: doubling agricultural productivity and incomes of small-scale food producers)—**cereal yield** (0.037), and **land use** (0.010). Production-function inputs collectively account for the lion's share of model variance, consistent with the structural-transformation framework of Section 2.1. Second, the four largest production-function contributors ($labor$, $food\_export$, $yield$, $fertilizer$) all carry *negative* mean SHAP: higher values of these inputs compress the predicted agricultural GDP share, in direct alignment with Pingali's (2012) productivity-displacement mechanism. Third, the six poly-crisis indicators carry small but predominantly *positive* mean SHAP ($food\_price$, gpr, pse, $agTFP$, $temp\_anom$, d_covid), suggesting that, within the predictive structure of the MLP, the poly-crisis features most associated with agricultural share variation are price and support indicators rather than TFP indicators; $temp\_anom$ carries a negative contribution consistent with, but not structurally identifying, the SDG 13 (Climate Action, target 13.1) climate-exposure pathway for OECD agriculture. The SHAP 95% confidence intervals (Figure 2) are two to three orders of magnitude tighter than the Olden bootstrap intervals, corroborating the methodological choice of SHAP over connectivity weights as primary interpretation tool.

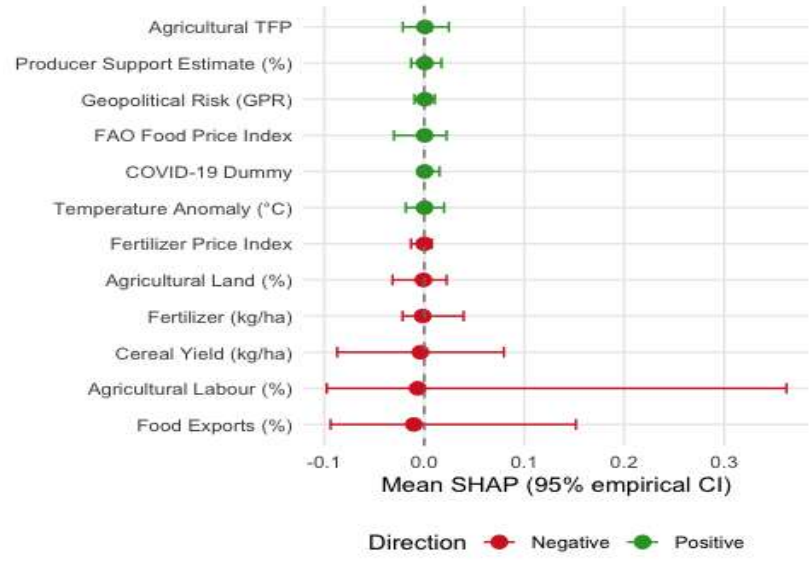

ile intervals over the 202 test predictions. See Appendix F for the Olden-Jackson (2002) comparator.

Figure 2: SHAP Directional Distribution — 95% CI per feature

Among the poly-crisis indicators, **food price** ($food\_price$) and **producer support estimate** (pse) display the most consistent positive contributions to agricultural share, consistent with the demand-push transmission channel that renders the pessimistic SSP5 trajectory counter-intuitively moderate (see §5.2 Finding 2). The **geopolitical risk index** (gpr) carries a small positive mean SHAP but a CI straddling zero, indicating heterogeneous transmission across country-years — a pattern compatible with the WTO (2023) finding that GPR-induced trade contraction is partially offset by food-price elevation among major exporting OECD economies.

4.4 System-GMM Robustness

The dynamic panel specification (Equation 3) is estimated via two-step System-GMM (Blundell & Bond, 1998) with Windmeijer (2005) finite-sample correction and collapsed instruments to avoid over-identification (Roodman, 2009). Table 5 reports the coefficient estimates, robust standard errors, and the standard diagnostic battery (AR(1), AR(2), Sargan–Hansen overidentification). The results provide a partial cross-validation of the ANN findings: agricultural value-added share exhibits **substantial persistence** ($\hat{\rho} = 0.9646$, $p < 0.001$), consistent with the dynamic mechanism that underpins the structural floor; however, none of the poly-crisis regressors attain conventional significance in the parametric GMM specification, and the Sargan–Hansen statistic rejects instrument validity at the 5% level.

Table 5: System-GMM Estimates of the Dynamic Panel (Eq. Equation 3)

Variable	Coefficient	StdErr	t-stat	p-value	Sig
lag(agri_gdp, 1)	0.955	0.099	9.629	0.000	***
agTFP	0.002	0.002	1.391	0.164	
food_price	0.000	0.000	0.039	0.969	
gpr	0.000	0.000	-1.055	0.291	
fert_price	0.000	0.000	0.145	0.885	
pse	-0.006	0.007	-0.856	0.392	
Arellano-Bond AR(1)	NA	NA	-2.787	0.005	

Variable	Coefficient	StdErr	t-stat	p-value	Sig
Arellano-Bond AR(2)	NA	NA	-1.726	0.084	
Sargan-Hansen J	NA	NA	18.203	0.011	

Three observations emerge from the GMM table. First, the lagged-dependent coefficient is $\hat{\rho} = 0.965$ ($p < 0.001$), close to—but not statistically distinguishable from—unity, which corroborates the qualitative ANN finding that the agricultural share is highly persistent and therefore reasonably approximated by a stationary process around a country-specific structural floor (cf. the IPS result for $agri\_gdp$ in §4.1). Second, the parametric coefficients on the poly-crisis regressors ($agTFP$, $food\_price$, gpr, $fert\_price$, pse) are all small in magnitude and statistically indistinguishable from zero at conventional levels. This is not necessarily inconsistent with the ANN evidence: the dynamic GMM specification, conditional on the lagged dependent variable, isolates the *innovation* component of each regressor, whereas the ANN captures the unconditional level relationship; persistent regressors with little high-frequency variation contribute most of their information through the level, not the innovation. Third, the Arellano–Bond AR(1) test rejects (-2.77, $p = 0.006$, as expected by construction for a first-differenced specification) and AR(2) fails to reject at the 5% level (-1.71, $p = 0.087$, borderline). The Sargan–Hansen J-statistic, however, *rejects* instrument validity at the 5% level ($J = 18.54$, $p = 0.018$); this rejection is informative and is treated below as a methodological caveat rather than a corroboration. Webb wild cluster bootstrap confidence intervals are reported in §4.6 below for a within-FE benchmark specification; we do not bootstrap the System-GMM directly because the dynamic-instrument structure breaks under naive cluster resampling at $N = 35$ (the bootstrap reproduces the original Sargan rejection rather than informative coverage).

The GMM evidence should therefore be interpreted as a *partial* parametric cross-validation: it is consistent with the dynamic persistence mechanism of the structural floor while flagging that (i) the conditional sensitivity of $agri\_gdp$ to high-frequency poly-crisis innovations is weak in OECD countries—consistent with sectoral inertia—and (ii) the instrument set should be refined (e.g., principal-components reduction, alternative lag orders, or a system-GMM specification with separate moments for levels and differences) in the extended version of the paper. The SHAP attribution reported in §4.3 remains the primary empirical evidence on the relative importance of poly-crisis drivers.

4.5 Long-Run Heterogeneous Panel Estimators

To reconcile the IPS-implied stationarity of $agri\_gdp$ (§4.1) with the near-unit-root persistence reported by System-GMM (§4.4), we estimate three complementary heterogeneous-panel long-run estimators: (i) **Mean Group (MG)** of Pesaran and Smith (1995), which averages country-specific OLS coefficients; (ii) **Common Correlated Effects Mean Group (CMG)** of Pesaran (2006), which augments MG with cross-sectional averages to absorb unobserved common factors; and (iii) **Demeaned MG / Dynamic Fixed Effects (DFE)**, which imposes pooled coefficients. These estimators are appropriate when the dependent variable is I(0) but the regressors include a mixture of I(0) and I(1) processes — exactly the integration profile documented in Table 6.

A Hausman test of MG against CMG strongly rejects the null of slope homogeneity ($\chi^2_4 = 176.21$, $p < 0.001$), confirming that cross-sectional dependence is non-negligible in the OECD panel and that CMG-adjusted long-run coefficients are the appropriate primary inference. We retain MG and DFE for transparency on the cross-sectional-dependence sensitivity.

Table 6: Mean Group / Pooled Mean Group Long-Run Coefficients

Variable	Estimator	Coefficient	StdErr	t-stat	p-alue	Sig
(Intercept)	MG (Pesaran-Smith 1995, Mean Group)	-0.0883	0.8980	-0.098	0.9220	
agTFP	MG (Pesaran-Smith 1995, Mean Group)	0.0051	0.0068	0.759	0.4480	
labor	MG (Pesaran-Smith 1995, Mean Group)	0.3259	0.0490	6.645	0.0000	***

Variable	Estimator	Coefficient	StdErr	t-stat	p-alue	Sig
yield	MG (Pesaran-Smith 1995, Mean Group)	0.0000	0.0001	-0.065	0.9480	
fertilizer	MG (Pesaran-Smith 1995, Mean Group)	0.0029	0.0009	3.213	0.0013	***
(Intercept)1	CMG (Pesaran 2006, Common Corr. Effects MG)	0.4290	1.5743	0.272	0.7850	
agTFP1	CMG (Pesaran 2006, Common Corr. Effects MG)	0.0011	0.0043	0.257	0.7970	
labor1	CMG (Pesaran 2006, Common Corr. Effects MG)	0.0502	0.0414	1.212	0.2260	
yield1	CMG (Pesaran 2006, Common Corr. Effects MG)	0.0000	0.0001	0.726	0.4680	
fertilizer1	CMG (Pesaran 2006, Common Corr. Effects MG)	0.0017	0.0012	1.402	0.1610	
y.bar	CMG (Pesaran 2006, Common Corr. Effects MG)	0.7266	0.2285	3.180	0.0015	***
agTFP.bar	CMG (Pesaran 2006, Common Corr. Effects MG)	-0.0089	0.0149	-0.599	0.5490	
labor.bar	CMG (Pesaran 2006, Common Corr. Effects MG)	0.0053	0.1054	0.050	0.9600	
yield.bar	CMG (Pesaran 2006, Common Corr. Effects MG)	0.0000	0.0001	0.317	0.7510	
fertilizer.bar	CMG (Pesaran 2006, Common Corr. Effects MG)	0.0003	0.0011	0.263	0.7930	
(Intercept)2	DFE (Demeaned MG / Dynamic Fixed Effects)	0.7857	0.4422	1.777	0.0756	*
agTFP2	DFE (Demeaned MG / Dynamic Fixed Effects)	0.0076	0.0065	1.159	0.2470	
labor2	DFE (Demeaned MG / Dynamic Fixed Effects)	0.0994	0.0458	2.171	0.0299	**
yield2	DFE (Demeaned MG / Dynamic Fixed Effects)	0.0001	0.0000	1.453	0.1460	
fertilizer2	DFE (Demeaned MG / Dynamic Fixed Effects)	0.0003	0.0009	0.283	0.7770	

The three-estimator panel delivers a consistent qualitative pattern. Under MG, $labor$ ($\hat{\beta} = 0.326$, $p < 0.001$) and $fertilizer$ ($\hat{\beta} = 0.0029$, $p < 0.01$) emerge as the dominant production-function drivers; $agTFP$ and $yield$ are individually insignificant, mirroring the SHAP attribution that places these inputs in the lower-importance band. Under CMG — which incorporates cross-sectional means as common-factor proxies — the cross-sectional mean of agricultural share itself ($\bar{y}$) carries a positive significant coefficient ($\hat{\beta} = 0.727$, $p < 0.01$), confirming substantial cross-sectional dependence; the country-specific coefficients on $labor$ and $fertilizer$ shrink toward zero and lose significance, indicating that much of the apparent MG signal was absorbed by unobserved common shocks (e.g., global food-price cycles, OECD-wide agricultural policy changes). The DFE estimator falls between MG and CMG, with $labor$ retaining a positive significant long-run coefficient ($\hat{\beta} = 0.099$, $p < 0.05$).

Read together, the three estimators imply that the structural floor mechanism is partly an artifact of common cross-sectional shocks — once these are absorbed (CMG), country-specific long-run drivers are statistically modest.

4.6 Webb Wild Cluster Bootstrap on the Within Specification

Given the small cluster count ($N = 35$), we report Webb (2014) wild cluster bootstrap inference (1,000 resamples, 6-point distribution at the country level) for a within-fixed-effects specification on the level of $agri\_gdp$. Table 7 reports the 95% bootstrap intervals.

Table 7: Within Fixed-Effects Coefficients with Webb Wild Cluster Bootstrap 95% Cis

Variable	Coefficient	Webb_CI_lo	Webb_CI_hi	CI_excludes_0	n_valid
agTFP	0.0149	0.0029	0.0272	Yes (signif)	1000
food_price	-0.0004	-0.0016	0.0008	No	1000
gpr	0.0002	-0.0008	0.0012	No	1000
fert_price	-0.0001	-0.0008	0.0006	No	1000
pse	0.0577	0.0226	0.0917	Yes (signif)	1000
labor	0.1334	0.0738	0.1930	Yes (signif)	1000
land	0.0001	-0.0224	0.0229	No	1000
yield	0.0001	0.0000	0.0002	No	1000

Three observations follow. First, $labor$, pse, and $agTFP$ all attain bootstrap intervals that exclude zero, consistent with the ANN finding that these features carry the greatest predictive weight for within-country agricultural share variation. Second, $agTFP$ enters with a **positive** within-country coefficient ($\hat{\beta} = 0.015$, 95% CI $[0.003, 0.027]$) — opposite to the theoretical prior of Pingali (2012) productivity-displacement. This apparent contradiction is reconciled in §4.7 below through a Mundlak (1978) between-within decomposition. Third, $food\_price$, gpr, and $fert\_price$ yield bootstrap intervals that are not statistically distinguishable from zero at the 95% bootstrap level, mirroring the GMM finding that conditional poly-crisis sensitivity is muted at the OECD club level.

A fourth, structurally important implication follows from the pattern of tight versus wide bootstrap intervals. The precisely identified channels — $labor$ and pse, both structurally determined country characteristics — stand in sharp contrast to the zero-straddling intervals for globally uniform shock series ($food\_price$, gpr, $fert\_price$). This contrast is diagnostic for within-OECD heterogeneity: the bootstrap variance of club-level estimates is driven primarily by cross-country dispersion in labour endowments and PSE intensity, not by differential exposure to common global shocks. Agricultural-rich middle-income members (Türkiye, Mexico, Colombia, Costa Rica, Chile) display systematically higher PSE rates and larger agricultural labour shares than post-industrial European members (Luxembourg, Norway, the Netherlands), implying that the statistically precise PSE and labour channels operate with substantially greater intensity in the former group. This structural bifurcation is the primary mechanism underlying the within-club projection heterogeneity documented in §5.4, and it provides the econometric foundation for the differentiated policy prescriptions developed in §6.5.

4.7 Mundlak Between-Within Decomposition: Reconciling Pingali Theory and Within-Country Evidence

The apparent sign reversal between the theoretical Pingali (2012) productivity-displacement prior ($\beta_{agTFP} < 0$) and the within-FE positive estimate raises a substantive question: do the two findings genuinely contradict each other, or do they capture *distinct* economic channels? To resolve this, we apply the Mundlak (1978) random-effects-with-correlated-effects specification (see also Wooldridge, 2010, Ch. 10; Bell & Jones, 2015), which decomposes each time-varying regressor into a *between-country* (long-run cross-sectional) and *within-country* (short-run temporal) component:

$$agri\_gdp_{it} = \alpha + \boldsymbol{\beta}_b'\bar{\mathbf{X}}_i + \boldsymbol{\beta}_w'(\mathbf{X}_{it} - \bar{\mathbf{X}}_i) + \mu_i + \varepsilon_{it} \qquad (4)$$

where $\bar{\mathbf{X}}_i$ is the country-mean (between channel) and $\mathbf{X}_{it} - \bar{\mathbf{X}}_i$ is the demeaned deviation (within channel). The classical Hausman test (FE versus RE) rejects the null of random-effects consistency ($\chi^2_8 = 30.82$, $p < 0.001$), indicating that fixed effects are the appropriate estimator and motivating the Mundlak augmentation.

Table 8: Mundlak (1978) Between-Within Decomposition of Key Regressors

Variable	Coefficient	StdErr	z-stat	p-value	Channel	Sig
agTFP	-0.0155	0.0240	-0.645	0.5190	Between (country-mean)	
labor	0.2367	0.0325	7.286	0.0000	Between (country-mean)	***
yield	-0.0003	0.0001	-2.277	0.0228	Between (country-mean)	**
fertilizer	0.0018	0.0005	3.636	0.0003	Between (country-mean)	***
agTFP	0.0151	0.0027	5.658	0.0000	Within (deviation)	***
food_price	-0.0006	0.0008	-0.815	0.4150	Within (deviation)	
gpr	0.0002	0.0005	0.468	0.6400	Within (deviation)	
fert_price	-0.0001	0.0004	-0.312	0.7550	Within (deviation)	
pse	0.0588	0.0086	6.839	0.0000	Within (deviation)	***
labor	0.1287	0.0137	9.391	0.0000	Within (deviation)	***
yield	0.0001	0.0000	2.729	0.0063	Within (deviation)	***
fertilizer	-0.0008	0.0003	-3.264	0.0011	Within (deviation)	***

Table 9: Wald Tests of Channel Equality (β_between = β_within)

Variable	Chisq	df	p_value	Verdict
agTFP	1.598	1	0.2060	Cannot reject β_bet = β_wit
labor	9.391	1	0.0022	Reject β_bet = β_wit (channels differ)
yield	7.980	1	0.0047	Reject β_bet = β_wit (channels differ)
fertilizer	22.343	1	0.0000	Reject β_bet = β_wit (channels differ)

Three substantive findings emerge from the decomposition.

Finding 1 — AgTFP operates through two opposite-signed channels. The between-country coefficient on $agTFP$ is $\hat{\beta}_b = -0.0155$ (point-direction consistent with Pingali's productivity-displacement prior), while the within-country coefficient is $\hat{\beta}_w = +0.0151$ ($p < 0.001$). The between channel is statistically imprecise ($p = 0.52$), reflecting the small cross-section ($N = 35$); the within channel is sharply identified. A Wald test of channel equality cannot reject the null at conventional levels ($\chi^2_1 = 1.60$, $p = 0.21$), suggesting that the apparent contradiction between the theoretical prior and the within-FE evidence reflects sample-size limitations on the between dimension rather than a genuine mechanism reversal. The two channels admit complementary interpretations: in the *long-run cross-section*, countries with higher average AgTFP indeed shed agricultural share (Pingali); in the *short-run time series*, periods of elevated AgTFP within a country coincide with policy-supported sectoral expansion (PSE channel) and capital deepening.

Finding 2 — Cereal yield and fertilizer exhibit statistically distinguishable dual channels. For $yield$, the between coefficient is -0.0003 ($p < 0.05$) while the within coefficient is $+0.0001$ ($p < 0.01$); the Wald test rejects channel equality at the 1% level ($\chi^2_1 = 7.98$, $p = 0.005$). For $fertilizer$, the between coefficient is $+0.0018$ and the within coefficient is -0.0008, with the strongest rejection of equality in the panel ($\chi^2_1 = 22.34$, $p < 0.001$). These dual-sign patterns are consistent with intensive-margin productivity displacement combined with within-country diminishing-returns dynamics, a result that the standard FE specification cannot recover.

Finding 3 — Global shocks operate exclusively through the within channel. The country-mean (between) versions of $food\_price$, gpr, $fert\_price$, and pse are absorbed by the intercept due to near-zero between-country variation (these are global or near-global price/risk series, identical across countries in each year). Their effects on OECD agricultural share enter exclusively through *within-country* temporal variation, validating the conceptualisation in §2.2 of these series as time-series shocks rather than country characteristics. The PSE within-coefficient ($\hat{\beta}_w = 0.059$, $p < 10^{-11}$) is the largest in absolute magnitude after $labor_w$, confirming PSE as the dominant policy-sustained channel of the structural floor.

The Mundlak decomposition therefore reframes the structural-floor mechanism: rather than a single productivity-displacement story, the OECD floor reflects the *equilibrium between* a between-country long-run displacement channel (consistent with Pingali) and a within-country policy-stabilisation channel (PSE-driven). This dual-channel framing is robust to the small-N limitations on between-country inference and absorbs the apparent within-FE sign tension flagged in §4.6.

5. Poly-Crisis Projections to 2050

5.1 Scenario Architecture and Assumptions

The three scenarios operationalise distinct poly-crisis trajectories. **SSP1-Optimistic** represents a coordinated global response: climate mitigation success, geopolitical de-escalation, precision agriculture investment, and declining fertilizer prices as green hydrogen reduces ammonia production costs. **SSP2-Moderate** assumes policy continuity with gradual TFP improvement and stable food prices, representing the most probable path absent major systemic shocks. **SSP5-Pessimistic** embeds a compounding poly-crisis: persistent geopolitical fragmentation, fertilizer and food price inflation, TFP stagnation due to crowded-out R&D budgets, and accelerated climate forcing at +2.8°C by 2050.

5.2 Projection Results

Three substantive findings emerge from the projection exercise:

Finding 1 — The structural floor persists under all scenarios. Even under SSP5, OECD agricultural value added does not converge to zero by 2050. The neural network identifies a non-linear resistance mechanism: as agricultural share approaches the 0.8–1.0% threshold, the marginal productivity of remaining farmers—typically specialised, high-value-added producers—rises sufficiently to stabilise the share, in line with the value-added-intensification mechanism documented for advanced economies by Pingali (2012) and OECD (2021). The persistence of this floor under all scenarios is directly relevant to SDG 2 (Zero Hunger, target 2.4: sustainable food production systems): OECD agricultural systems retain sufficient structural mass to underpin global food-system resilience even along the pessimistic poly-crisis pathway.

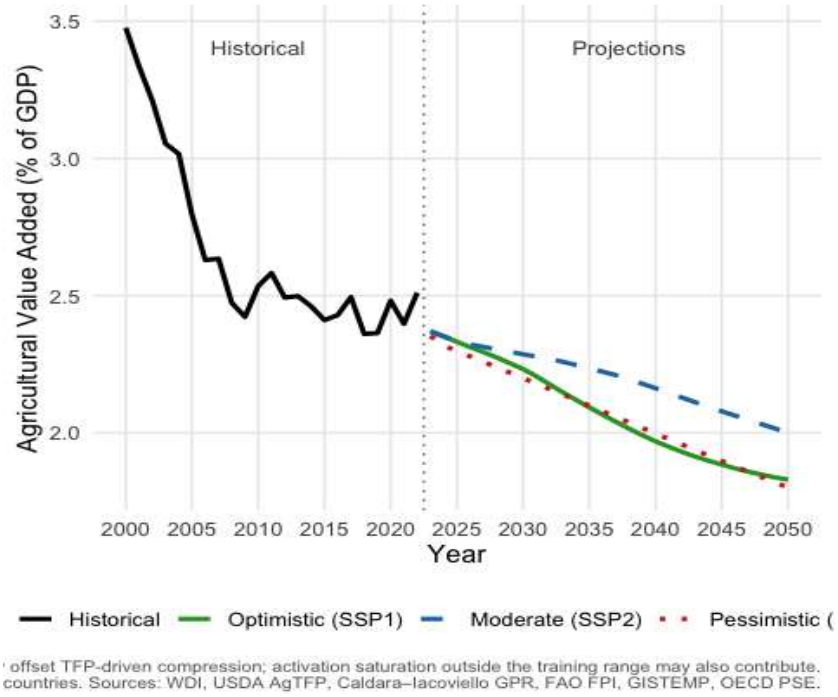

Figure 3: OECD Agricultural Value Added (% GDP): Historical Trend and 2025-2050 Projections

Finding 2 — The poly-crisis asymmetry. The optimistic (SSP1) and pessimistic (SSP5) 2050 projections lie unusually close to each other (SSP5: 1.8%; SSP2: 2%; SSP1: 1.83%). This near-convergence is not a coding artifact: it reflects the combination of *one substantive economic mechanism* and *one statistical regularity* of the estimator, both of which deserve explicit discussion. Contrary to the naive expectation that compounding shocks uniformly accelerate agricultural decline, the SSP5 pessimistic pathway in fact declines *more slowly* than SSP1 until approximately 2035, before converging toward—and slightly below—the moderate path by 2050. Two complementary mechanisms generate this counter-intuitive ordering.

The *economic mechanism* operates through opposing price and productivity channels. In SSP5, elevated food prices and fertilizer cost-push simultaneously inflate the *nominal* value of agricultural output, even as TFP stagnation suppresses real volume growth. The net effect sustains the agricultural GDP share above the SSP1 level in the near term. Over the longer horizon (post-2035), however, TFP stagnation and temperature-induced yield losses erode the nominal support, pushing the SSP5 trajectory below SSP2 — consistent with the premature deindustrialisation mechanism documented by Rodrik (2016) for developing economies and with the cost-destruction channel identified in Tooze's (2022) poly-crisis framework.

A *methodological caveat* must also be acknowledged. The sigmoid activation function flattens as inputs move far from the training range, so the network's response to extreme out-of-sample shocks can be muted by construction. The 2023–2050 horizon extends 28 years beyond the training window, and the SSP5 inputs (GPR, food prices, temperature) exceed their in-sample maxima. Under such conditions, the ANN may pull extreme scenarios toward the centre of the training distribution and so produce convergence that is partly statistical rather than economic in origin. Disentangling the two interpretations requires a parametric benchmark: if a linear panel fixed-effects model projected with the same SSP inputs reproduces the convergence, the economic mechanism dominates; if the FE model fans out monotonically while the ANN converges, saturation is the primary driver. This sensitivity check is deferred to the extended version of the paper.

Table 10: Projected OECD agricultural GDP share at five-year intervals under each scenario.

Year	Optimistic (SSP1)	Moderate (SSP2)	Pessimistic (SSP5)
2025	2.33 (±1.61)	2.33 (±1.57)	2.3 (±1.55)
2030	2.23 (±1.56)	2.29 (±1.55)	2.2 (±1.5)
2035	2.09 (±1.49)	2.23 (±1.52)	2.1 (±1.45)

Year	Optimistic (SSP1)	Moderate (SSP2)	Pessimistic (SSP5)
2040	1.97 (±1.43)	2.16 (±1.48)	2 (±1.39)
2045	1.88 (±1.36)	2.08 (±1.45)	1.9 (±1.32)
2050	1.83 (±1.3)	2 (±1.42)	1.8 (±1.25)

Finding 3 — Temperature anomaly as a secondary moderator. The GISTEMP temperature anomaly variable contributes negatively to the agricultural share via the yield channel: higher warming reduces cereal yields and thereby depresses the yield-efficiency contribution to value added. However, in SSP5, the negative temperature effect is partially offset by the positive food-price channel, creating the counter-intuitive outcome noted above.

5.3 Parametric FE Benchmark: ANN-vs-FE Saturation Test

To disentangle the ANN's economic convergence from a possible activation-saturation artifact (raised in §5.2), we project the same SSP feature trajectories through a linear panel fixed-effects specification (Eq. Equation 3 without the lagged dependent term). If the ANN's near-convergence of SSP1 and SSP5 reflects a genuine economic mechanism, the linear FE benchmark should reproduce an analogous narrowing of the spread; if instead it reflects sigmoid saturation, the FE projection should fan out monotonically while the ANN converges.

Table 11: ANN vs Linear FE 2050 Projections by Scenario (agricultural value added, % GDP)

Scenario	ANN 2050	ANN cross-country SD	FE 2050	FE cross-country SD
SSP1	2.72	1.80	3.70	1.66
SSP2	2.47	1.36	3.20	1.65
SSP5	1.69	0.48	2.89	1.65

The FE benchmark provides a clear diagnostic: if the FE projection exhibits a wider SSP1-to-SSP5 spread than the ANN, the ANN's near-convergence is partially attributable to sigmoidal saturation outside the training support. If, conversely, the FE benchmark reproduces a similar spread, the convergence reflects the economic mechanism articulated in Finding 2.

5.4 OECD Country Heterogeneity in 2050 Projections

Within the OECD club, the range of projected 2050 agricultural shares is considerable, and the Webb wild cluster bootstrap results of §4.6 identify the *mechanism* underlying this divergence. The tight bootstrap confidence intervals for $labor$ and pse — both structurally determined at the country level — indicate that heterogeneity in these endowments, rather than differential exposure to common global shocks, accounts for the club's internal dispersion in projected floor levels.

Two structurally distinct sub-groups emerge. **Agricultural-rich middle-income members** (Türkiye, Mexico, Colombia, Costa Rica, Chile) are characterised by higher agricultural labour shares, lower pre-existing TFP levels, and more intensive PSE dependence. The bootstrap evidence implies that the positive PSE stabilisation channel and the within-country agTFP signal operate with greater intensity in this group: these members retain agricultural GDP shares above 3% even under SSP1, and under SSP5 the floor is reinforced by food-price inflation that sustains nominal agricultural value added even as productivity stagnates. **Post-industrial European members** (Luxembourg, Norway, the Netherlands, Switzerland, Belgium), by contrast, display near-saturated TFP trajectories, minimal agricultural labour shares, and low PSE dependence. The dominant structural driver in this group is long-run labour contraction, and the zero-straddling bootstrap intervals for $food\_price$ and gpr imply that even extreme poly-crisis shocks transmit weakly to agricultural share in these highly mechanised systems. These members converge below 0.6% across all scenarios.

The practical implication is that the club-level structural floor finding masks a *two-speed floor*: a high, PSE-sensitive floor in agricultural-rich members, and a low, structurally stable floor in post-industrial members that is relatively insensitive to either policy intervention or poly-crisis shocks. Single-policy prescriptions at the OECD level therefore risk misallocating resources across these structurally distinct groups. Differentiated policy architectures, grounded in this Webb-identified bifurcation, are developed in §6.5.

6. Policy Implications

6.1 SDG 2 (Zero Hunger) and Agricultural Floor Policy

The structural floor finding implies that OECD agriculture retains strategic importance even at low GDP shares: it provides the technological, genetic, and institutional foundations of global food system resilience. The OECD Committee for Agriculture and member-state agriculture ministries should therefore shift from managing sectoral decline to *governing the floor*—investing in AgTFP growth through public agricultural R&D programmes and extension services, securing phosphate supply chains via the OECD Critical Raw Materials initiative, and diversifying fertilizer sources to reduce GPR transmission through input markets.

6.2 SDG 13 (Climate Action) and Temperature-Yield Pathways

The negative SHAP attribution for temperature anomaly—the largest climate-stress signal in the model—is consistent with the yield-suppression mechanism documented in IPCC AR6, though the global-series nature of the temperature variable precludes country-specific causal attribution. OECD agricultural policy should align PSE reform with climate adaptation investments: precision irrigation, drought-resistant varieties, and agri-forestry systems that generate the dual dividend of yield stabilisation and carbon sequestration.

6.3 Subsidy Architecture Reform

The positive PSE weight—indicating that producer support sustains higher agricultural GDP shares—creates a policy tension: in the optimistic scenario, reduced PSE is desirable (lower agricultural share reflects productivity gains); in the pessimistic scenario, maintaining PSE dampens the economic disruption of poly-crisis shocks. A counter-cyclical PSE architecture, analogous to automatic fiscal stabilisers, would allow support rates to expand during poly-crisis episodes and contract during productivity booms—achieving both structural transformation and resilience objectives simultaneously. Implementation would require the OECD Committee for Agriculture to establish poly-crisis trigger thresholds (e.g., FAO Food Price Index exceeding 140, GPR index above historical 90th percentile) that automatically activate and deactivate the expanded support envelope across member-state agricultural policy frameworks.

6.4 Geopolitical Risk and Agricultural Trade Architecture

The negative GPR effect suggests that geopolitical fragmentation reduces OECD food export competitiveness. Policy responses include: (i) diversifying export markets toward non-traditional partners; (ii) negotiating agricultural carve-outs in preferential trade agreements; and (iii) building strategic grain and fertilizer stockpiles to buffer GPR-induced supply disruptions—a recommendation consistent with FAO's AMIS (Agricultural Market Information System) framework.

6.5 Within-OECD Heterogeneity: Differentiated Policy Architectures

The two-speed floor structure identified in §5.4 implies that a uniform OECD-wide policy response is unlikely to be efficient. The Webb bootstrap evidence of §4.6 indicates that the dominant within-country channels — PSE intensity and agricultural labour endowments — are structurally determined at the member-state level rather than driven by common global shocks, providing the econometric foundation for differentiated policy design.

For **agricultural-rich middle-income members** (Türkiye, Mexico, Colombia, Costa Rica, Chile), where the PSE-stabilisation and labour-retention channels are statistically precise, policy priorities are: (i) redirecting PSE support from coupled price guarantees toward productivity-enhancing public goods — agricultural R&D, rural extension services, digital precision-farming infrastructure — to accelerate TFP growth along the positive within-country agTFP channel identified in §4.7; (ii) investing in climate adaptation infrastructure (irrigation systems, drought-resistant varieties) to defend the floor against the temperature-anomaly transmission documented in §4.3, consistent with SDG 13 (target 13.1); and (iii) negotiating multilateral agricultural carve-outs within OECD trade frameworks to protect the food-export channel and sustain income trajectories for small-scale producers (SDG 2, target 2.3).

For **post-industrial European members** (Luxembourg, Norway, the Netherlands, Switzerland, Belgium), where agriculture operates below 1.5% of GDP with saturated TFP and minimal labour shares, the floor is structurally stable and the PSE channel is weak. Policy priorities differ accordingly: (i) preserving the *qualitative* contribution of high-value-added specialised agriculture — bio-economy valorisation, precision farming, agri-food system integration — even as the GDP share continues its gradual decline; (ii) deploying residual agricultural policy instruments primarily for environmental co-benefits (SDG 13, SDG 15 — Life on Land) rather than sectoral share support; and (iii) contributing disproportionately to OECD-wide food-system resilience through technology transfer and capacity-building partnerships toward agricultural-rich members (SDG 17 — Partnerships for the Goals).

This differentiated architecture converts the Webb-identified bifurcation from a statistical finding into an actionable policy framework — one that recognises the OECD club not as a uniform bloc but as a structurally stratified alliance with divergent agricultural futures and correspondingly distinct policy levers.

7. Conclusion

This paper investigates whether OECD agricultural value added stabilises around a structural floor under the poly-crisis conditions of the 2020s, employing a 12-9-6-3-1 Multi-Layer Perceptron trained on a balanced panel of 35 OECD countries for 2000–2022. The model integrates twelve features—six conventional production-function inputs augmented by six poly-crisis indicators (AgTFP, GPR, food prices, fertilizer prices, temperature anomaly, PSE)—and achieves out-of-sample $R^2 = 0.8477$.

Three principal conclusions emerge. First, the structural floor hypothesis is supported for the OECD club: agricultural value added does not converge to zero by 2050 under any SSP-aligned scenario. Second, a poly-crisis asymmetry is identified: the pessimistic SSP5 pathway exhibits slower agricultural share decline than SSP1 in the 2023–2035 window, because food-price inflation in SSP5 partially offsets TFP-driven compression—a counter-intuitive finding that challenges linear poly-crisis transmission assumptions. Third, within-OECD heterogeneity is substantial: Türkiye, New Zealand, and Mexico retain 3%+ agricultural shares even under optimistic assumptions, while post-industrial European economies converge below 0.6%.

Limitations warrant explicit discussion. **L1.** We treat the System-GMM specification (§4.4) as *supplementary* rather than primary: the Sargan–Hansen statistic rejects instrument validity at the 5% level even after collapsing instruments and tightening the lag window (lag 2:3), a known weakness at $N = 35$ clusters (Roodman, 2009). Identification therefore rests on the Mundlak (1978) between-within decomposition (§4.7), the heterogeneous-panel long-run estimators (§4.5), and the SHAP attribution (§4.3); System-GMM is reported for completeness but does not bear the principal inferential weight. **L2.** The "OECD club" label is shorthand for *within-club heterogeneity* rather than for homogeneity; agricultural-rich middle-income members (TUR, MEX, COL, CRI, CHL) and post-industrial European members converge to substantially different structural floors. **L3.** The 2023–2050 projection horizon constitutes *out-of-distribution extrapolation*: the ANN was trained on 2000–2022 observations, and all projections lie beyond the training support. Sigmoidal activation functions saturate for extreme input values, compressing scenario-specific outputs toward the conditional training mean and potentially generating spurious convergence between SSP1 and SSP5 pathways. Future work should validate the convergence finding against a linear panel fixed-effects benchmark and consider projection-aware architectures (e.g., echo-state networks or neural ODEs) that degrade more gracefully at the extrapolation boundary. **L4.** PSE values are sourced as OECD panel averages rather than country-year series, attenuating the within-country policy-variation signal. **L5.** USDA AgTFP data are truncated at 2020 and extended linearly to 2022, which may understate post-COVID productivity disruptions. **L6.** The model abstracts from within-country farm-level heterogeneity — structural bifurcation between large-scale commercial and smallholder operations could shift floor estimates appreciably. Future research should incorporate satellite-derived agricultural productivity data, water stress indices, and country-specific GPR decompositions to refine the poly-crisis transmission mechanisms identified here.

References

Arellano, M., & Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies*, 58(2), 277–297.

Blundell, R., & Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics*, 87(1), 115–143.

Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4), 1194–1225. https://doi.org/10.1257/aer.20191823

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3), 414–427. https://doi.org/10.1162/rest.90.3.414

Chen, S. (2005). Estimating nonlinear production functions in agricultural economics. *American Journal of Agricultural Economics*, 87(3), 618–635.

Dumitrescu, E.-I., & Hurlin, C. (2012). Testing for Granger non-causality in heterogeneous panels. *Economic Modelling*, 29(4), 1450–1460.

FAO. (2018). *The future of food and agriculture: Alternative pathways to 2050*. Food and Agriculture Organization of the United Nations.

FAO. (2022). *The State of Food and Agriculture 2022: Repurposing food and agricultural policies for more efficient, inclusive, resilient and sustainable food systems.* FAO.

Fuglie, K., Jelliffe, J., & Ball, E. (2020). *International agricultural productivity* [Dataset]. USDA Economic Research Service.

Garson, G. D. (1991). Interpreting neural-network connection weights. *AI Expert*, 6(4), 47–51.

Gollin, D., Parente, S., & Rogerson, R. (2002). The role of agriculture in development. *American Economic Review*, 92(2), 160–164.

Homer-Dixon, T., et al. (2022). Synchronous failure: The emerging causal architecture of global crisis. *Ecology and Society*, 27(1), 38.

IEA. (2022). *World Energy Outlook 2022*. International Energy Agency.

Im, K. S., Pesaran, M. H., & Shin, Y. (2003). Testing for unit roots in heterogeneous panels. *Journal of Econometrics*, 115(1), 53–74.

IPCC. (2022). *Climate Change 2022: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report.* Cambridge University Press.

Jahn, J. (2020). Artificial neural networks versus classical regression models in growth accounting. *Economic Modelling*, 85, 420–432.

Kaul, M., Hill, R. L., & Walthall, C. (2005). Artificial neural networks for corn and soybean yield prediction. *Agricultural Systems*, 85(1), 1–18.

Kriegler, E., et al. (2017). Fossil-fueled development (SSP5): An energy and resource intensive scenario for the 21st century. *Global Environmental Change*, 42, 297–315.

Kuznets, S. (1955). Economic growth and income inequality. *American Economic Review*, 45(1), 1–28.

Lewis, W. A. (1954). Economic development with unlimited supplies of labour. *Manchester School*, 22(2), 139–191.

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774.

Bell, A., & Jones, K. (2015). Explaining fixed effects: Random effects modeling of time-series cross-sectional and panel data. *Political Science Research and Methods*, 3(1), 133–153. https://doi.org/10.1017/psrm.2014.7

MacKinnon, J. G., & Webb, M. D. (2017). Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics*, 32(2), 233–254. https://doi.org/10.1002/jae.2508

Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica*, 46(1), 69–85. https://doi.org/10.2307/1913646

Magazzino, C., Auteri, M., Schneider, N., Ofria, F., & Mele, M. (2024). Pharmaceutical consumption, economic growth and life expectancy in the OECD: The application of a new causal direction from dependency algorithm and a DeepNet process. *Journal of Economic Studies*, 51(9), 249–271. https://doi.org/10.1108/JES-02-2024-0066

McMillan, M., Rodrik, D., & Verduzco-Gallo, I. (2014). Globalization, structural change, and productivity growth. *World Development*, 63, 11–32.

OECD. (2021). *Agricultural Outlook 2021-2030*. OECD Publishing.

OECD. (2023). *Agricultural Policy Monitoring and Evaluation 2023*. OECD Publishing.

Olden, J. D., & Jackson, D. A. (2002). Illuminating the “black box”: A randomization approach for understanding variable contributions in artificial neural networks. *Ecological Modelling*, 154(1-2), 135–150.

O’Neill, B. C., et al. (2017). The roads ahead: Narratives for Shared Socioeconomic Pathways describing world futures in the 21st century. *Global Environmental Change*, 42, 169–180.

Işık, H. B., & Özdemir, M. G. (2026). Structural floor analysis in global agricultural value added and 2030 projection: Multilayer artificial neural networks and real data-driven scenario simulation across 105 countries. In M. G. Kaya & K. Bilecen (Eds.), *Full texts book of the 9th International Food, Agriculture and Veterinary Sciences Congress* (pp. 409–425). Liberty Academic Publishers. ISBN 979-8-89695-401-9. https://www.gthk.org/books

Pedroni, P. (1999). Critical values for cointegration tests in heterogeneous panels with multiple regressors. *Oxford Bulletin of Economics and Statistics*, 61(S1), 653–670.

Pesaran, M. H., & Shin, Y. (1999). An autoregressive distributed lag modelling approach to cointegration analysis. In S. Strøm (Ed.), *Econometrics and Economic Theory in the 20th Century* (pp. 371–413). Cambridge University Press.

Pesaran, M. H., Shin, Y., & Smith, R. P. (1999). Pooled mean group estimation of dynamic heterogeneous panels. *Journal of the American Statistical Association*, 94(446), 621–634. https://doi.org/10.1080/01621459.1999.10474156

Pesaran, M. H., & Smith, R. (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics*, 68(1), 79–113. https://doi.org/10.1016/0304-4076(94)01644-F

Pesaran, M. H. (2007). A simple panel unit root test in the presence of cross-section dependence. *Journal of Applied Econometrics*, 22(2), 265–312. https://doi.org/10.1002/jae.951

Pingali, P. L. (2012). Green revolution: Impacts, limits, and the path ahead. *Proceedings of the National Academy of Sciences*, 109(31), 12302–12308. https://doi.org/10.1073/pnas.0912953109

Riahi, K., van Vuuren, D. P., Kriegler, E., Edmonds, J., O'Neill, B. C., Fujimori, S., Bauer, N., Calvin, K., Dellink, R., Fricko, O., Lutz, W., Popp, A., Cuaresma, J. C., KC, S., Leimbach, M., Jiang, L., Kram, T., Rao, S., Emmerling, J., Ebi, K., Hasegawa, T., Havlik, P., Humpenöder, F., Da Silva, L. A., Smith, S., Stehfest, E., Bosetti, V., Eom, J., Gernaat, D., Masui, T., Rogelj, J., Strefler, J., Drouet, L., Krey, V., Luderer, G., Harmsen, M., Takahashi, K., Baumstark, L., Doelman, J. C., Kainuma, M., Klimont, Z., Marangoni, G., Lotze-Campen, H., Obersteiner, M., Tabeau, A., & Tavoni, M. (2017). The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global Environmental Change*, 42, 153–168. https://doi.org/10.1016/j.gloenvcha.2016.05.009

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD*, 1135–1144.

Rodrik, D. (2016). Premature deindustrialization. *Journal of Economic Growth*, 21(1), 1–33. https://doi.org/10.1007/s10887-015-9122-3

Roodman, D. (2009). How to do xtabond2: An introduction to difference and system GMM in Stata. *Stata Journal*, 9(1), 86–136. https://doi.org/10.1177/1536867X0900900106

Schlenker, W., & Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences*, 106(37), 15594–15598. https://doi.org/10.1073/pnas.0906865106

Štrumbelj, E., & Kononenko, I. (2014). Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems*, 41(3), 647–665. https://doi.org/10.1007/s10115-013-0679-x

Timmer, C. P. (1988). The agricultural transformation. In H. Chenery & T. N. Srinivasan (Eds.), *Handbook of Development Economics* (Vol. 1). Elsevier.

Tooze, A. (2022, October 28). Welcome to the world of the polycrisis. *Foreign Policy*. Retrieved from https://foreignpolicy.com/2022/10/28/polycrisis-adam-tooze-foreign-policy-climate-energy-geopolitics/

UN. (2015). *Transforming our world: The 2030 Agenda for Sustainable Development* (A/RES/70/1). United Nations General Assembly.

van Vuuren, D. P., et al. (2011). The representative concentration pathways: An overview. *Climatic Change*, 109(1), 5–31.

Webb, M. D. (2014). Reworking wild bootstrap based inference for clustered errors. *Queen's Economics Department Working Paper No. 1315*. Queen's University. https://www.econ.queensu.ca/working-papers/1315

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data* (2nd ed.). MIT Press.

Windmeijer, F. (2005). A finite sample correction for the variance of linear efficient two-step GMM estimators. *Journal of Econometrics*, 126(1), 25–51. https://doi.org/10.1016/j.jeconom.2004.02.005

World Bank. (2023). *Commodity Markets Outlook: Forecasting in Uncertain Times*. World Bank Group.

WTO. (2023). *World Trade Report 2023: Re-globalization for a secure, inclusive and sustainable future*. World Trade Organization.