WORKING PAPER 05/2026

Green Public Spending Shocks: Composition Effects and State-Dependent Heterogeneity Across Countries

João Tovar Jalles John Beirne Donghyun Park

The South East Asian Central Banks (SEACEN) Research and Training Centre (80416-M)

Kuala Lumpur, Malaysia

WORKING PAPER 05/2026

Green Public Spending Shocks: Composition Effects and State-Dependent Heterogeneity Across Countries

João Tovar Jalles John Beirne Donghyun Park

September 2026

Table of Contents

List of Figures

List of Tables

Appendix

Green Public Spending Shocks: Composition Effects and State-Dependent Heterogeneity Across Countries[1]

João Tovar Jalles[2] John Beirne[3] Donghyun Park[4]

Abstract

We quantify “green fiscal multipliers” by constructing country-specific, residual-based shocks to environmental protection outlays (COFOG) and tracing their effects on macroeconomic and environmental outcomes in a multi-country panel. Reduced-form estimates imply that an unexpected 1-pp-of-GDP increase in environmental spending lowers contemporaneous real GDP per-capita growth by about 2.6 pp globally, IV estimates are larger in magnitude (≈–4 pp globally; ≈–7 pp in EMDEs). Local projections show a small short-run dip that troughs around year 2 at –0.05% and fades thereafter, with mild medium-run rebound. Heterogeneity is pronounced: EMDEs drive the near-term contraction; effects are more negative at high public debt and in surplus/open economies and are mitigated where adaptive capacity (ND-GAIN readiness) is high. The findings imply that green fiscal policy should be carefully sequenced and tailored to country conditions to manage short-term costs while maximising long-term benefits. For central banks, green fiscal expansions should be incorporated into the assessment of the effectiveness and transmission of monetary policy. Accounting for their scale, timing, and country-specific conditions can help to minimise the risk of policy miscalibration and support price and financial stability during the transition to a low-carbon economy.

Keywords: Environmental protection spending; green fiscal multipliers; residual-based shocks; local projections; state dependence; EMDE vs AE

JEL Codes: E62; Q58; C23; E32; O44

---

[1] We gratefully acknowledge the valuable feedback and suggestions provided participants at an ADB-APAEA workshop on “Climate change, debt management, and fiscal policy design”, held in Bangkok in June 2026. The views and opinions expressed in this paper are solely those of the authors and do not necessarily reflect those of their respective institutions or employers. The authors declare that they have no relevant financial or personal relationships that could appear to influence the research presented in this paper. Any remaining errors are the sole responsibility of the authors.
[2] Corresponding author. University of Lisbon-Lisbon School of Economics and Management (ISEG), Rua do Quelhas 6, 1200-781 Lisboa, Portugal. Research in Economics and Mathematics (REM) and Research Unit on Complexity and Economics (UECE), Universidade de Lisboa-ISEG, Rua Miguel Lupi 20, 1249-078 Lisbon, Portugal. Economics for Policy, Universidade Nova de Lisboa-Nova School of Business and Economics, Rua da Holanda 1, 2775-405 Carcavelos, Portugal. IPAG Business School, 184 Boulevard Saint-Germain, 75006 Paris, France. Email: joaojalles@gmail.com
[3] Asian Development Bank, 6 ADB Avenue, Mandaluyong City 1550, Metro Manila, Philippines. Email: jbeirne@adb.org
[4] The SEACEN Centre, Kuala Lumpur Malaysia. Email: donghyun.park@seacen.org

1. Introduction

Public budgets increasingly finance "green" goods—waste and wastewater systems, pollution control, biodiversity and ecosystem services, and small but growing envelopes for environmental R&D. Yet policy debates often talk past a central macroeconomic tension: green outlays can carry front-loaded costs while benefits arrive with a lag. Compliance mandates and infrastructure build-outs tend to be capital-intensive, import-heavy, and administratively complex; productivity and resilience gains materialise only after equipment is installed, networks are connected, and firms and households adapt. This paper quantifies those near-term trade-offs and the conditions under which they are attenuated or amplified.

Two simple observations motivate our analysis. First, environmental protection (COFOG 05) is small in levels but volatile at the margin. Median spending is roughly 0.55–0.70% of GDP in advanced economies (AEs) and 0.25–0.35% in EMDEs, with a wide interquartile range; a one-percentage-point-of-GDP change is therefore large relative to typical budget shares. Second, the composition is uneven: waste and wastewater dominate everywhere, pollution abatement rises with stringency (especially in AEs), and environmental R&D remains modest. These facts hint that composition and country circumstances should matter for macro effects.

Anecdotally, the mechanisms are intuitive. Municipal wastewater upgrades or SO_2/NO_x abatement programmes require imported machinery and specialised contractors; execution lags and co-financing can crowd out private activity in the year of implementation. By contrast, once plants are operating and compliance technologies disseminate, firms can reoptimise, energy intensity falls, and pollution exposure declines. Our contribution is to discipline this narrative with a global panel design that separates exogenous innovations in environmental spending from endogenous responses to the cycle, and that traces dynamic adjustment paths.

Our approach proceeds in three steps. First, we construct country-specific, residual-based shocks to environmental protection outlays. For each country, spending (as % of GDP) is regressed on its own lag and lagged fundamentals (log income per capita, growth, public debt, age dependency); the first-differenced residual, rescaled so that one unit equals 1 pp of GDP, is our innovation (winsorised at the 1st–99th percentiles). This procedure purges trend and persistence while leaving unanticipated, discretionary changes. Second, we estimate reduced-form panels with country and year fixed effects and robust (Driscoll–Kraay) inference, and complementary 2SLS specifications that use lagged shocks as instruments for observed spending changes. Third, we turn to local projections (LPs) to recover impulse responses over eight years, including state-dependent LPs that allow responses to vary with fiscal space (public debt), external exposure (trade openness), and adaptive capacity. For the latter, we use ND-GAIN measures: *Vulnerability* captures exposure/sensitivity across critical systems, while *Readiness* (labelled "resilience" in figures for readability) explicitly embeds a governance block (regulatory quality, rule of law, control of corruption) alongside economic and social capacity; higher readiness is therefore a practical proxy for stronger institutions and implementation capacity relevant to executing green projects.

The data set is a large, unbalanced panel of AEs and EMDEs from 1990 to 2025, combining government spending by COFOG 05 and subcomponents (pollution abatement, environmental R&D, waste management, wastewater) with standard macro controls and environmental outcomes (emissions per capita, energy intensity, renewable electricity share). We document stylised facts on levels, dispersion, and composition, then bring the shock series to the reduced-form, IV, and LP frameworks.

Three sets of results emerge. (i) Contemporaneous panels. Globally, an unexpected 1-pp-of-GDP increase in environmental spending is associated with about –2.6 percentage points lower per-capita growth on impact (5% significance); IV estimates are more negative (≈ –4 pp globally, ≈ –7 pp in EMDEs). Subcomponents point to pollution abatement as the clearest driver of near-term contraction; waste and wastewater are negative but less precisely estimated; environmental R&D shows a large but imprecise drag, consistent with lumpy programmes and deferred payoffs. (ii) Dynamics. The global LP shows a small trough around year 2 (≈ –0.05%), statistically different from zero through year 2, followed by a gradual recovery that crosses zero around year 5 and becomes mildly positive thereafter. (iii) Heterogeneity. EMDEs drive the global near-term contraction; AEs are negative but statistically indistinct from zero. Conditioning reveals strong state dependence: high debt amplifies the short-run decline; very open economies display attenuated net effects (import leakage and external adjustment cushion the cycle); and higher ND-GAIN readiness does not remove the initial dip but shortens its duration and raises medium-run payoffs. Robustness checks—panel Granger non-causality, predictability regressions, pre-trend event studies, placebos, and orthogonality to spending levels—support the shocks' exogeneity and timing.

The remainder of the paper is organised as follows. Section 2 reviews the related literature. Section 3 lays out a simple theoretical framework that clarifies channels from environmental spending to output and environmental outcomes. Section 4 describes data and methods. Section 5 presents and discusses the results. Section 6 concludes with policy implications.

2. Literature Review

2.1 Theoretical Framework

Theoretical work on fiscal multipliers has traditionally emphasised the role of government spending in stimulating demand and supporting capital accumulation (e.g., Barro, 1990; Baxter and King, 1993). In these models, public expenditure raises aggregate demand in the short-run and may augment productive capacity in the long-run through infrastructure investment. However, such frameworks generally treat the environment as an externality or a constraint rather than as a productive input.

More recent contributions in the field of environmental macroeconomics integrate environmental quality into the production function and explicitly model the interplay between growth, damages, and policy intervention (e.g., Acemoglu et al., 2012; Nordhaus, 2014; Heutel, 2012). In these settings, public spending directed toward adaptation, mitigation, or protection generates dual benefits: (i) a conventional Keynesian demand effect and (ii) a supply-side effect through improved environmental resilience and reduced damages. The magnitude of the fiscal multiplier is therefore state-dependent: it may be larger in economies highly exposed to climate risks or with low baseline environmental quality.

A further strand of the literature examines the intertemporal trade-off between traditional and green public spending. On the one hand, diverting resources toward environmental protection may reduce short-run output if such spending is less investment-intensive or has weaker demand spillovers (Pereira and Pereira, 2017). On the other hand, by improving long-run productivity and lowering the probability of climate-induced losses, environmental investment can yield higher multipliers over medium horizons (Batini et al., 2021). These theoretical insights suggest that the effectiveness of environmental public spending depends critically on structural conditions—such as vulnerability, fiscal space, and institutional quality—underscoring the relevance of the heterogeneity analysis we pursue in this paper.

2.2 Empirical Evidence

Empirical research on fiscal multipliers has grown extensively, but the subset focusing on "green" multipliers remains limited. The conventional literature estimates government spending multipliers using cross-country panels or local projections, often finding values between 0.5 and 1.5 depending on the state of the economy, the type of expenditure, and fiscal conditions (Auerbach and Gorodnichenko, 2012; Ramey and Zubairy, 2018). Infrastructure spending has typically been shown to yield higher multipliers because of its capital-augmenting role (Bom and Ligthart, 2014).

Within environmental economics, studies have begun to quantify the macroeconomic effects of climate adaptation and mitigation policies. For example, Hepburn et al. (2020) show that green investment programmes following the global financial crisis produced employment and growth benefits comparable to or greater than traditional stimulus. Batini et al. (2021) estimate that public investment in climate adaptation yields larger long-run multipliers in climate-vulnerable economies. Similarly, IMF (2022) evidence suggests that green spending not only stimulates output but also enhances resilience to climate shocks, leading to more stable growth paths.

At the same time, empirical evidence also points to important sources of heterogeneity. Multipliers are larger in economies with spare capacity, accommodative monetary policy, or strong institutions (Ilzetzki et al., 2013). For green multipliers specifically, vulnerability indices such as ND-GAIN have been shown to condition effectiveness, with more exposed countries deriving greater benefits from environmental spending (Hallegatte et al., 2016). Moreover, fiscal space shapes outcomes: high-debt economies may experience weaker multipliers as credibility concerns offset stimulus effects (Auerbach and Gorodnichenko, 2017).

Taken together, the existing literature suggests that environmental public spending can generate sizable macroeconomic returns, particularly in countries facing elevated climate risks. However, systematic cross-country evidence on the magnitude and dynamics of these “green fiscal multipliers” remains scarce. This paper seeks to fill that gap by combining a panel reduced-form approach with local projections, exploiting novel residual-based shocks to environmental spending, and explicitly testing for heterogeneity by income level, fiscal stance, and climate vulnerability.

3. A Simple Theoretical Framework

We develop a stylised model to clarify the channels through which environmental public spending affects both economic activity and environmental outcomes. The framework is intentionally simple, designed to motivate our empirical specifications rather than to provide a fully structural account.

Output in country i at time t depends on physical capital $K_{i,t}$, effective labour, $L_{i,t}$, and the quality of the environment $Q_{i,t}$:

$$Y_{i,t} = A_{i,t} K_{i,t}^{\alpha} L_{i,t}^{1-\alpha} Q_{i,t}^{\eta}, \qquad 0 < \alpha < 1, \eta > 0 \tag{1}$$

where $A_{i,t}$ is total factor productivity. The exponent η captures how environmental quality enhances productivity by reducing health costs, mitigating climate damages, and improving resilience. Since our empirical analysis focuses on real GDP per capita growth, the relevant object is Y/L, where the contribution of labour growth is netted out.

The government collects revenues $T_{i,t}$ (e.g., general taxes) and allocates them across two types of expenditure:

- *Traditional public spending ($G_{i,t}^{T}$),* such as infrastructure, defence, or health, which augments the capital stock.
- *Environmental protection spending ($G_{i,t}^{E}$),* such as waste management, biodiversity protection, or environmental R&D, which improves environmental quality.

$$T_{i,t} = G_{i,t}^{T} + G_{i,t}^{E} \tag{2}$$

Environmental quality evolves as:

$$Q_{i,t+1} = (1 - \delta_Q) Q_{i,t} + \theta G_{i,t}^{E} - \chi Y_{i,t} \tag{3}$$

where δ_Q is the natural depreciation of environmental assets, θ measures the effectiveness of environmental spending in raising quality, and χ captures degradation from economic activity (emissions, pollution). Importantly, θ and χ are heterogeneous across countries: vulnerable economies face higher χ, while resilient economies enjoy higher θ. Empirically, we capture this heterogeneity by conditioning on the ND-GAIN vulnerability and resilience indices in the local projections.

Physical capital evolves as:

$$K_{i,t+1} = (1 - \delta_K)K_{i,t} + I_{i,t} + \phi G_{i,t}^T \quad (4)$$

where δ_K is the depreciation rate, $I_{i,t}$ denotes private investment, and ϕ captures the efficiency of traditional public spending in augmenting the capital stock.

In the short-run, reallocating funds toward $G_{i,t}^E$ may crowd out traditional investment and yield smaller immediate demand effects. Over time, however, higher environmental spending increases $Q_{i,t}$, which:

1. Enhances productivity (η>0).
2. Reduces effective vulnerability to climate and environmental shocks (lower χ).
3. Stabilises long-run growth paths, particularly in climate-exposed economies.

This framework motivates our empirical strategy and generates several testable predictions:

- **Positive multipliers**: Environmental spending should raise both output per capita and environmental quality.
- **Heterogeneity by vulnerability and resilience**: Effects are larger where χ is high (greater damages per unit of output) and where θ is low (spending less effective), i.e., in countries with high ND-GAIN vulnerability and low ND-GAIN resilience.
- **Role of fiscal space**: Crowding out of traditional investment or high debt service may dampen short-run effects.
- **Dynamic gains dominate static trade-offs**: Growth effects are expected to be stronger in the medium run, as productivity and resilience benefits accumulate, compared to the immediate period.

4. Methodology and Data

This section outlines our empirical strategy and data sources. A key innovation relative to prior work is the construction of country-specific residual-based shocks to environmental public spending, which address concerns of endogeneity and policy cyclicality. We then employ both reduced-form panel regressions and dynamic local projections to estimate the effects of these shocks on macroeconomic and environmental outcomes.

4.1 Construction of Residual-Based Environmental Spending Shocks

To capture plausibly exogenous changes in environmental public expenditure, we proceed in two steps. First, we estimate a separate regression for each country *i,* relating environmental protection spending (as a share of GDP) to its own persistence and a set of lagged domestic fundamentals:

$$EP_{i,t} = \alpha_i + \phi_i(L)EP_{i,t-1} + \beta_i^{'}X_{i,t-1} + u_{i,t} \quad (5)$$

where $EP_{i,t}$ denotes environmental protection spending as a share of GDP; α_i captures country fixed effects; $\phi_i(L)$ is a lag polynomial to absorb persistence; $X_{i,t-1}$ is a vector of lagged domestic fundamentals: log real GDP per capita, real GDP growth, public debt-to-GDP, and the age-dependency ratio. The regression is estimated separately for each country, allowing all coefficients to vary across national contexts. No year fixed effects or global controls are included at this stage, so that international co-movements are not differenced out mechanically.).

The resulting residuals $u_{i,t}$ capture the component of environmental spending that is orthogonal to these domestic fundamentals and to its own persistence. In the second step, we take first differences to remove low-frequency drift and then rescale the residual innovations to units of percentage points of GDP:

$$Shocki_{i,t}^{EP} = \Delta u_{i,t}. \quad (6)$$

This measure represents unexpected innovations in environmental public spending that are not systematically linked to contemporaneous macroeconomic conditions. This measure captures unexpected innovations in environmental spending that are not systematically related to contemporaneous macroeconomic conditions. By construction, these shocks isolate the unpredictable fiscal component of green spending decisions.

We construct shocks both for total environmental protection expenditure and for its four main COFOG subcomponents: i) Pollution abatement — spending on reducing air, soil, and water pollution through regulation, monitoring, and direct abatement programmes; ii) Environmental R&D — expenditures on research, development, and technology programmes with an environmental objective; iii) Waste management — spending on collection, treatment, and disposal of solid waste; iv) Wastewater management — expenditures on sewage networks, treatment plants, and related infrastructure.

Each subcomponent is treated analogously: for every country, we residualise the series against its own lag, plus the same set of domestic controls, and construct first-differenced residuals scaled into percentage points of GDP. These orthogonalised shocks reflect exogenous innovations in the composition of environmental spending, enabling us to explore whether the macroeconomic and environmental effects differ across categories.

In the second-stage panel regressions, we relate macroeconomic and environmental outcomes to these shocks, including country and year fixed effects and a parsimonious set of lagged controls (real income, trade openness, and debt). Variables used in the first-stage construction are not reintroduced in the second-stage equations, avoiding redundancy and ensuring that the estimated effects are not driven by mechanical re-inclusion of the conditioning set.

4.2 Reduced-form Panel Regressions

We begin by estimating reduced-form panel regressions of the form:

$$\Delta y_{i,t} = \theta\, Shocki_{i,t}^{EP} + \lambda\, W_{i,t-1} + \mu_i + \tau_t + \varepsilon_{i,t} \quad (7)$$

where $\Delta y_{i,t}$ denotes real GDP per-capita growth (constructed as 100×ΔlogGDPpc; in extensions we also consider environmental outcomes (e.g., emissions per capita). The vector $W_{i,t-1}$ collects lagged macro controls—log income per capita, public debt (% GDP), trade openness, inflation, private credit (% GDP), the terms of trade, and the current account balance (% GDP). Country fixed effects μ_i absorb time-invariant heterogeneity; year dummies τ_t control for common shocks. The regressor $Shocki_{i,t}^{EP}$ is the first-differenced residual from the country-by-country orthogonalisation described in Section 4.1, rescaled so that one unit equals one percentage point of GDP. To limit the influence of outliers, we winsorise this series at the 1st–99th percentiles, which is our baseline treatment variable in the panel.

We estimate the model with country and year fixed effects and Driscoll–Kraay standard errors to allow for heteroskedasticity, serial correlation, and cross-sectional dependence. In this setup, θ captures the contemporaneous association between an unanticipated one-pp-of-GDP increase in environmental spending and per-capita growth, conditional on lagged fundamentals and common shocks. This provides a transparent benchmark for the sign and order of magnitude of short-run effects before turning to dynamics.

Because the shocks are residual innovations, they can also serve as instruments for actual changes in environmental spending. Accordingly, we complement the reduced-form with 2SLS specifications instrumenting the alternative main regressor $\Delta EP_{i,t}$ with lagged values of the shock (e.g., $Shocki_{i,t-1}^{EP}$,$Shocki_{i,t-2}^{EP}$). IV regressions include the same FE and controls and use Driscoll–Kraay standard errors. We report first-stage diagnostics (Kleibergen–Paap F-statistics, Stock–Yogo critical values) and Anderson–Rubin confidence sets.

To study how short-run effects vary with initial conditions, we interact the shock with centered lagged states and plot marginal effects across the empirical range (p10–p90).[5]

Reduced-form estimates are inherently static and can understate delayed benefits of environmental outlays. Many programmes (waste management, biodiversity, pollution control) feature planning and execution lags, and standard fiscal multipliers propagate gradually through demand, investment, and labour markets. We therefore turn to local projections in the next section to trace the full adjustment path and separate short-run demand effects from medium-run productivity and resilience gains.

[5] For transparency, we build explicit lag copies, enforce same-sample masks for the interaction regressions, and compute confidence bands from the Driscoll–Kraay covariance matrix. These figures remain part of the reduced-form analysis because they quantify static state dependence at impact.

4.3 Dynamic Effects via Local Projections

We study how exogenous innovations in environmental spending propagate to macroeconomic and environmental outcomes using the local projections (LP) method of Jordà (2005). LPs are well suited to our setting: they estimate horizon-specific responses without imposing a full multivariate dynamic structure, which is valuable in a heterogeneous, multi-country panel with limited time length. Outcomes include real GDP per-capita growth (our baseline), private investment, and employment, as well as environmental indicators drawn from the WDI—CO_2 emissions per capita, energy intensity of GDP, renewable electricity share. This mix lets us speak to both the macroeconomic footprint and the environmental efficacy of green fiscal shocks. As detailed in Section 4.1, the spending innovation is the first-differenced residual from country-specific orthogonalisation, scaled so that one unit equals one percentage point of GDP and winsorised at the 1st–99th percentiles in the baseline. For each horizon h=1,…,8, we estimate the following LP:

$$\Delta y_{i,t+h} = \alpha_{i,h} + \delta_{t,h} + Shocki_{i,t}^{EP} + \sum_{c=0}^{1} \pi_{c,h} X_{i,t-c} + \sum_{m=0}^{h} \kappa_{m,h} Shocki_{i,t+m}^{EP} + u_{i,t+h} \qquad (8)$$

where $\Delta y_{i,t+h} = y_{i,t+h} - y_{i,t}$ is the h-period cumulative change in the dependent variable; $Shocki_{i,t}^{EP}$ is the innovation defined in Section 4.1; and $X_{i,t}$ includes inflation, public debt, private credit, trade openness, the current account, and the terms of trade. We include country fixed effects $\alpha_{i,h}$ and year–horizon fixed effects $\delta_{t,h}$. The lead terms $\{Shocki_{i,t+m}^{EP}\}_{m=1}^{h}$ implement the Teulings–Zubanov (2014) correction for overlapping horizons.[6]

To discipline dynamics, we conduct a grid search over lag orders for (i) the treatment shock *s*, (ii) the dependent variable *d*, and (iii) the controls *c* (each in {0,…,4}), minimising $AIC_{adj} = AIC + 2(s + d + c)$. The selected structure is s=0, d=0, c=1. Accordingly: Shock: enters contemporaneously only (L(0)); no shock lags are included; Anticipation: we include leads F(1..h) of the shock (Teulings–Zubanov (2014) correction); Outcome dynamics: no lags of Δy (since d=0); Controls: $X_{i,t}$ and $X_{i,t-1}$ only (because c=1). For environmental outcomes (e.g., CO_2 per capita, energy intensity), we additionally include the baseline level of the indicator to accommodate mean reversion and differences in trend. This parsimonious setup is not only data-driven; it also reduces the risk of over-fitting in a panel with many countries but moderate T. In robustness checks, adding one to three lags of Δy or of the shock leaves the IRFs within baseline confidence bands, supporting the chosen specification.

Estimation proceeds equation-by-equation for each horizon with Spatial Correlation Consistent (Driscoll–Kraay, 1998) standard errors, which are robust to heteroskedasticity, serial correlation, and cross-sectional dependence—a key concern in globally synchronised cycles and commodity shocks. As usual, the effective sample declines with hhh because later horizons require longer outcome windows. Fisher-type panel unit-root tests support stationarity of the growth-rate variables and weak persistence in levels once differenced;

[6] Including $\delta_{t,h}$ controls for cross-sectional dependence from common shocks but also conditions out the shared component of the innovation and its aggregate impact on activity. With $\delta_{t,h}$ in place, the IRFs reflect country-specific responses relative to the global trend at each horizon.

sums of dynamic coefficients remain small across horizons, suggesting that finite-sample distortions from short T are limited.

Macroeconomic variables come from the World Bank´s World Development Indicators.

4.4 Stylised Facts

Using COFOG data on general government environmental protection outlays as a share of GDP, Figure 1 summarises the cross-sectional median and interquartile range (IQR) annually for the global sample, advanced economies (AEs), and emerging market and developing economies (EMDEs). Four regularities emerge. First levels differ markedly by income group. AE medians are persistently around 0.55–0.70% of GDP, roughly double EMDE medians (≈ 0.25–0.35%). The global median sits between the two and tracks shifts in the EMDE share of the sample. This gap is consistent with higher fiscal capacity and stronger environmental mandates in AEs. Second, dispersion is sizable—and asymmetric. The IQR is wide in all groups, with the upper quartile in AEs approaching ~1% of GDP in several years, while the lower quartile in EMDEs often lies near 0.10–0.20%. This indicates substantial heterogeneity in policy effort and institutional implementation across countries at any point in time. Third, trends are flat to mildly declining globally; EMDEs stabilise at lower levels. Since the mid-2000s, the global median is broadly flat to slightly lower. In EMDEs, the median settles at low but gradually rising levels after early volatility around 2000; AEs show a relatively stable median with episodic upticks. Fourth and finally, cross-sectional heterogeneity dominates time variation. Year-to-year movements in the medians are modest compared with the width of the IQRs, suggesting strong policy inertia within countries and persistent differences across them. This motivates the use of country fixed effects and the exploration of state dependence (e.g., fiscal space, openness, adaptive capacity) in the empirical analysis.

These facts have two implications for the paper's empirical strategy. First, the wide IQRs justify splitting the sample by income group and conditioning responses on country fundamentals. Second, the mild aggregate trends and pronounced cross-sectional dispersion reinforce the value of local projections with year–horizon fixed effects, which identify country-specific responses relative to the common global component.

Figure 1. Environmental Protection Expenditure as a Share of GDP: Median and Interquartile Range Across Country Groups (1990–2023)

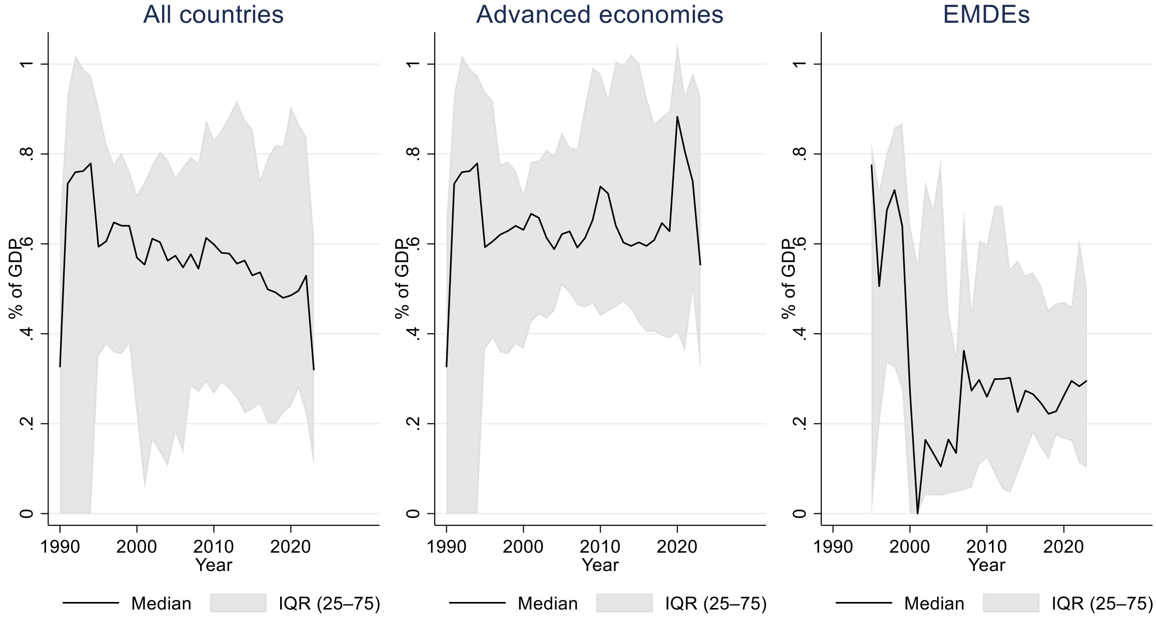

Note: Median is the solid black line. 25th and 75th percentiles denoted by the grey bands. Axes are common across panels.
Source: Authors using COFOG data.

Figure 2 shows how the composition of environmental protection outlays evolves across AEs and EMDEs at four dates (1990, 2000, 2010, 2020). In both groups, waste-related items dominate the envelope throughout, with wastewater taking the largest share in the early years—especially in EMDEs—consistent with basic network expansion and sanitation catch-up. Over time the centre of gravity shifts toward waste management, reflecting urbanisation and rising solid-waste needs, while wastewater's share recedes as networks mature. AEs allocate a visibly larger fraction to pollution abatement—peaking around 2010—and maintain a slightly higher, though still modest, commitment to environmental R&D; EMDEs devote a smaller and steadier share to pollution and similarly low R&D. By 2020 the two groups look more alike in the balance between waste management and wastewater, yet they do not converge in innovation intensity, with R&D remaining marginal in both. This composition helps interpret our impact estimates: pollution abatement is compliance-driven and capital-intensive, plausibly tied to larger short-run output drags, whereas waste and wastewater are infrastructure heavy with long gestation and higher import content, pointing to implementation lags and demand leakage rather than immediate stimulus.

Figure 2. Composition of Environmental Protection Expenditure by Function (Share of Total Environmental Protection), 1990–2020

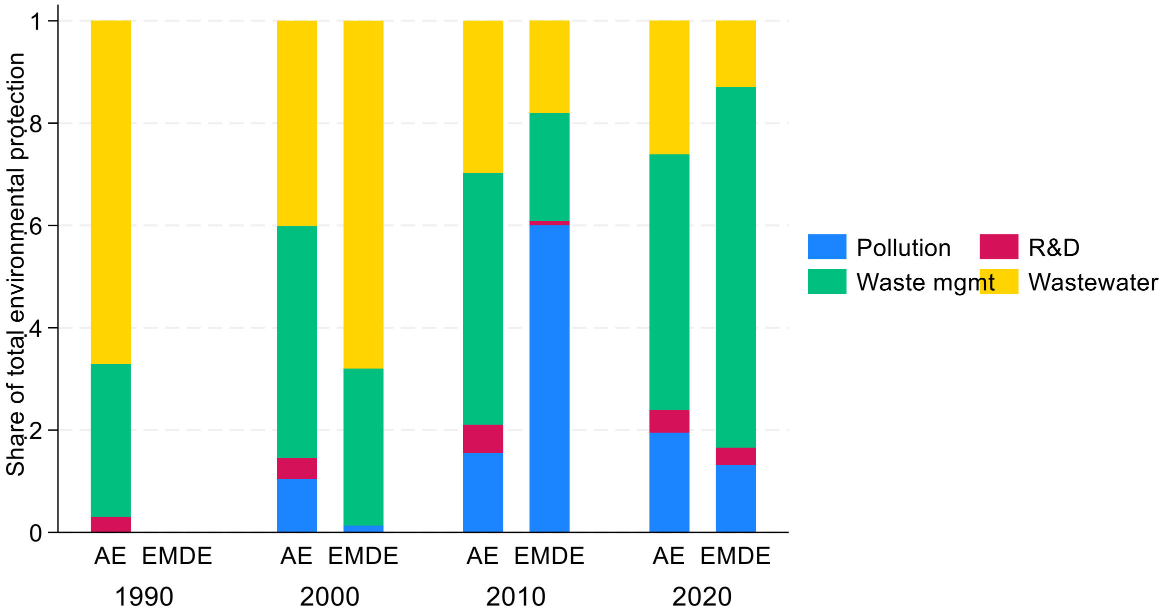

Note: Bars report the share of each component in total environmental protection expenditure. Blue = Military, Red = Civil, Green = R&D, Yellow = Foreign aid. Figures shown for Advanced Economies (AE) and Emerging Markets and Developing Economies (EMDEs) in benchmark years 1990, 2000, 2010, and 2020.
Source: Authors using IMF COFOG data.

In Appendix Figure A1, the distribution of environmental spending shocks (percentage points of GDP) is shown. The mass is tightly centred at zero and the kernel density exhibits thin tails, indicating that most realisations are small. This pattern is consistent with the construction of shocks as first-differenced, residual innovations that have been winsorised at the 1st–99th percentiles, so that infrequent, unanticipated adjustments are captured rather than trend movements. In Appendix Figure A2, decade-by-decade box–dot plots are presented. Medians remain essentially at zero across the 1990s, 2000s, 2010s, and 2020s, and interquartile ranges are relatively stable. Occasional tail observations are observed in all decades, with a mild thickening in the 2000s and 2020s, reinforcing the interpretation of a stationary, mean-zero innovation. In Appendix Figure A3, the time series for the United Kingdom, as an ad-hoc example, is displayed. A transitory positive spike is followed by a sharp negative reversal in the late 2000s, after which fluctuations are small and mean-reverting around zero. This profile is what would be expected when discretionary changes or one-off execution effects are superimposed on a persistent expenditure path. In Appendix Figure A4, the contemporaneous relationship between the shock and the level of environmental spending is depicted using a scatter with a lowess fit. A near-flat curve centred at the origin is obtained, indicating that the innovations are not mechanically tied to the spending level and that orthogonality is broadly achieved by the residualisation and differencing steps. Taken together, these panels provide evidence that the shock series is centred and stationary, largely orthogonal to contemporaneous spending, and suitable for tracing dynamic responses in the local-projection analysis.

A complete list of countries included in the sample, along with their classification as AEs or EMDEs, is provided in Appendix. The coverage of variables varies across countries due to data availability.

5. Empirical Results

5.1 Baseline Panel Results

We begin by relating residual innovations in environmental protection outlays—scaled so that one unit equals 1 pp of GDP—to annual real GDP per-capita growth. In the global panel, the reduced-form coefficient on the shock is –2.58 (s.e. 1.08), significant at 5%. Interpreted literally, an unanticipated 1 pp-of-GDP increase in environmental spending is associated with a 2.6 percentage-point reduction in contemporaneous per-capita growth. This short-run contraction is consistent with mechanisms emphasised in the fiscal-policy literature—front-loaded compliance and investment costs, execution lags, and crowding-out of private activity when capacity or financing constraints bind (Ramey, 2011; Ilzetzki, Mendoza and Végh, 2013; Auerbach and Gorodnichenko, 2012).

Subsample estimates for AEs and EMDEs are also negative (–2.22 and –1.67, respectively) but imprecisely estimated, which is unsurprising given the smaller samples (212 AE and 815 EMDE country-years) and wider standard errors. The direction of the effect is nevertheless common across groups, suggesting similar near-term dynamics. Among controls, the coefficients have the expected signs. Lagged income is negative and significant in the global and EMDE samples, indicating conditional convergence. Trade openness is positive and significant throughout: a 10-percentage-point increase in openness is associated with roughly 0.2–0.6 pp higher growth. Inflation is negatively correlated with growth, while terms-of-trade improvements are growth-enhancing. The current-account balance enters negatively in the global and EMDE samples, consistent with the idea that a contemporaneous move toward surplus often coincides with weaker domestic absorption. Debt and private credit are not robustly significant in the short-run once these controls and fixed effects are included.

Table 1. Baseline Reduced-form Panel Estimates: Growth Response to Environmental Spending Shocks

Specification Sample	(1) Global	(2) AE	(3) DEV
Environmental spending shock	-2.578**	-2.219	-1.667
	(1.083)	(1.718)	(2.240)
Log real GDP per capita (t-1)			
	-5.015***	-3.449	-5.354***
Public debt (% of GDP) (t-1)	(1.820)	(3.070)	(1.794)
	0.810	0.172	1.656
Trade openness (% of GDP) (t-1)	(1.150)	(0.800)	(1.307)
	0.040***	0.020**	0.062***
Consumer price inflation (annual %) (t-1)	(0.013)	(0.007)	(0.020)
	-0.050**	-0.033	-0.052**
Private sector credit (% of GDP) (t-1)	(0.019)	(0.055)	(0.020)
	-0.009	-0.014	0.020
Change in terms of trade (t-1)	(0.009)	(0.009)	(0.027)
	0.168***	-0.006	0.180***
Current-account balance (% of GDP) (t-1)	(0.047)	(0.148)	(0.057)
L.currentaccount	-0.075***	-0.093	-0.069**
	(0.026)	(0.068)	(0.033)
Obs.	1035	212	815

Note: The dependent variable is annual real GDP per-capita growth (percentage points). The key regressor is Environmental spending shock, a country-specific environmental-spending innovation constructed as the first-difference of residuals from country-by-country regressions of environmental protection outlays (share of GDP) on their own lag and lagged domestic fundamentals; the series is rescaled so that 1 unit = 1 percentage point of GDP and winsorised at the 1st–99th percentiles. Controls are lagged one year. All specifications include country fixed effects and year fixed effects; coefficients are therefore identified from within-country variation relative to the global co-movement in each year. Standard errors (in parentheses) are Spatial Correlation Consistent à la Driscoll–Kraay (1998). "Global," "AE," and "DEV" refer to the full sample, advanced economies, and EMDEs (developing/emerging), respectively; the panel is unbalanced, so coverage varies by country-year. Asterisks denote significance: * p<0.10, ** p<0.05, *** p<0.01.

To translate innovations into causal multipliers of observed spending changes, we estimate 2SLS specifications that instrument the shock with its first and second lags (Angrist and Pischke, 2009). The IV coefficients are more negative than the reduced-form: –4.00 globally (s.e. 1.98, $p<0.10$) and –7.22 in EMDEs (s.e. 2.32, $p<0.01$), while the AE point estimate is similar in sign but imprecise, reflecting limited power. The larger IV magnitudes suggest that reduced-form estimates may understate the contemporaneous causal effect—consistent with measurement error in the shock or partial offsetting adjustment within the period (Romer and Romer, 2010; Mertens and Ravn, 2012). First-stage diagnostics (Kleibergen–Paap statistics and Stock–Yogo critical values) indicate adequate instrument relevance in the global and EMDE samples; the AE first stage is weaker, in line with the small AE sample. Control coefficients remain stable across RF and IV, reinforcing the interpretation that the contractionary impact is tied to the exogenous green-spending innovation rather than to changes in macro conditions captured by the covariates.

Taken together, Tables 1–2 point to a short-run output cost of environmental outlays—larger in EMDEs—likely reflecting higher import content, tighter financing constraints, and absorptive-capacity frictions typical of developing economies (Ilzetzki et al., 2013; Kraay, 2014). These costs need not be permanent: theory and case evidence point to medium-run productivity and resilience gains once projects come online (Ramey, 2020; Jordà, 2005). We therefore turn to local projections to trace adjustment paths over several years and to quantify how effects vary with state variables such as fiscal space, openness, and climate vulnerability (Auerbach and Gorodnichenko, 2013; Teulings and Zubanov, 2014).

Table 2. Instrumental-variables (2SLS) Panel Estimates: Causal Effect of Environmental Spending Innovations on Growth

Specification	(1)	(2)	(3)
Sample	Global	AE	EMDE
Environmental spending shock	-4.004*	-4.198	-7.224***
	(1.982)	(3.530)	(2.323)
Log real GDP per capita (t-1)	-5.113**	-3.413	-5.523***
	(1.870)	(3.007)	(1.953)
Public debt (% of GDP) (t-1)	0.736	-0.058	1.591
	(1.125)	(1.030)	(1.270)
Trade openness (% of GDP) (t-1)	0.040***	0.020**	0.065***
	(0.014)	(0.009)	(0.021)
Consumer price inflation (annual %) (t-1)	-0.050**	-0.037	-0.051**
	(0.021)	(0.060)	(0.022)
Private sector credit (% of GDP) (t-1)	-0.011	-0.014	0.013
	(0.009)	(0.009)	(0.025)
Change in terms of trade (t-1)	0.168***	0.011	0.168**
	(0.055)	(0.139)	(0.063)
Current-account balance (% of GDP) (t-1)	-0.077**	-0.093	-0.078**
	(0.030)	(0.071)	(0.033)
Observations	1,035	212	815
R-squared	0.418	0.712	0.415
Country FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Note: The dependent variable is annual real GDP per-capita growth (percentage points). The potentially endogenous regressor is Environmental spending shock (defined as in Table 1), instrumented with its first and second lags. The control set, fixed effects, and variable definitions match Table 1. Results are obtained by two-stage least squares with country and year fixed effects. Standard errors are computed with a covariance estimator that is robust to heteroskedasticity, serial correlation, and cross-sectional dependence, using a three-year window and a finite-sample adjustment. The reported R2 refers to the second-stage regression. First-stage diagnostics (available on request) include the Kleibergen–Paap rk Wald F-statistic, Stock–Yogo weak-IV critical values, and Anderson–Rubin confidence sets, which are robust to weak instruments. “Global,” “AE,” and “EMDE” denote the full sample and income-group splits; the panel is unbalanced.
Standard errors are in parentheses; significance levels: * p<0.10, ** p<0.05, *** p<0.01.

Figure 3 shows how the impact of a one–percentage-point-of-GDP environmental spending shock on real GDP per-capita growth varies with initial conditions. Estimates come from the reduced-form panel with country and year fixed effects, lagged macro controls, and a covariance matrix robust to heteroskedasticity, serial correlation, and cross-sectional dependence; the line is the point estimate and the shaded area is the 95% confidence band evaluated over the p10–p90 range of each state variable.

Figure 3. Marginal Effect of a One–percentage-point-of-GDP Environmental Spending Shock on Real GDP per-capita growth, by Initial Conditions

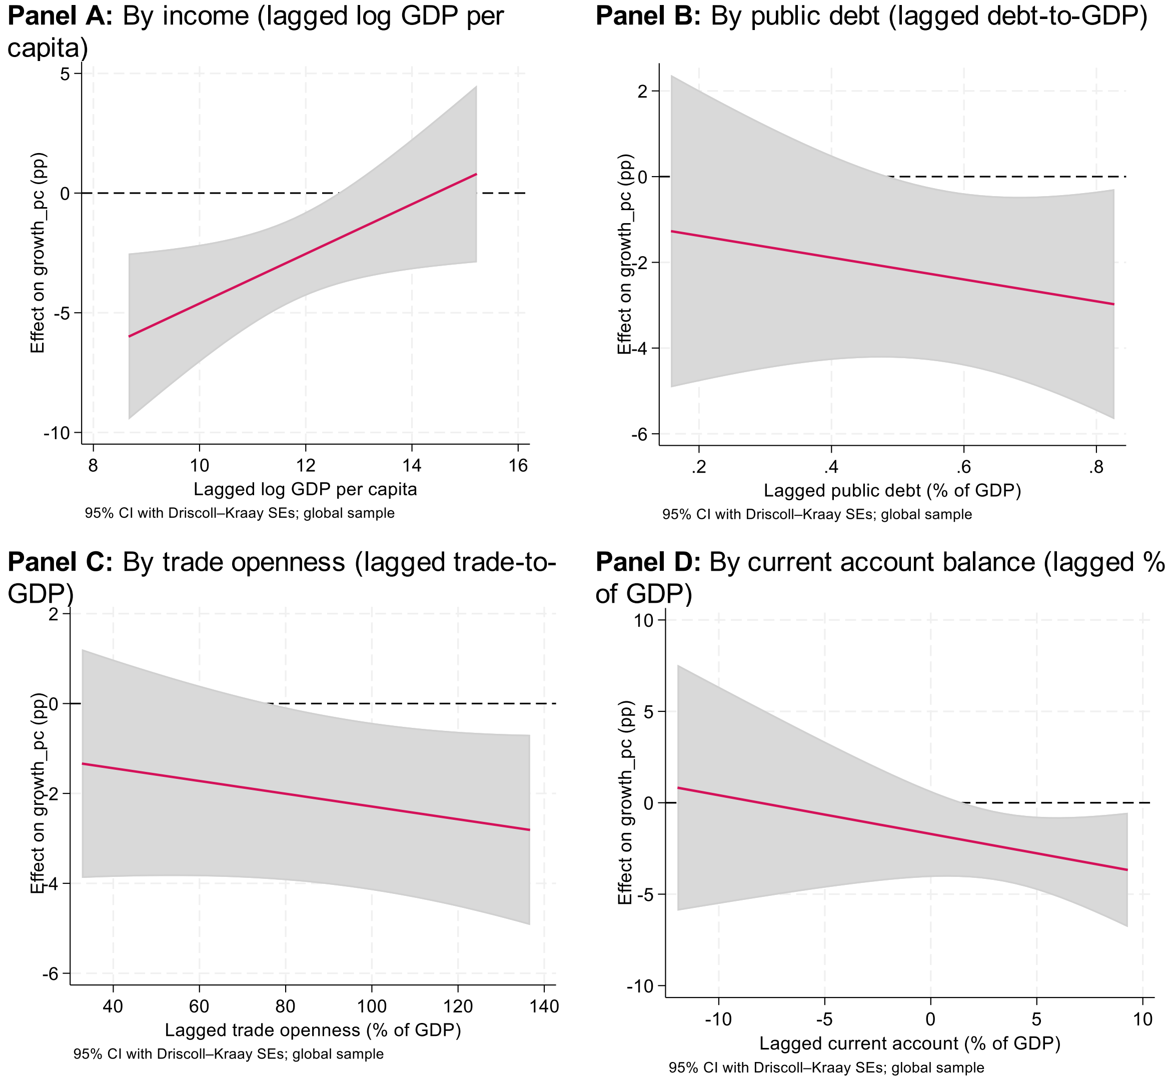

Note: Each panel reports the estimated marginal effect conditional on a lagged state variable. The solid line is the point estimate; the shaded band is the 95% confidence interval; the horizontal dashed line marks zero effect. Estimates come from the reduced-form panel with country and year fixed effects and lagged controls (lagged log income, public debt, trade openness, inflation, private credit, terms of trade, and current account). The shock is the first-differenced residual innovation in environmental protection outlays, scaled so that one unit equals 1 percentage point of GDP and winsorised at the 1st–99th percentiles. Confidence intervals are computed with a covariance estimator robust to heteroskedasticity, serial correlation, and cross-sectional dependence using a three-year window. Effects are plotted over the p10–p90 range of each state variable. Global sample; unbalanced panel.

Panel A conditions on lagged income and displays a clear income gradient: at low income levels, the effect is significantly negative—the confidence band lies below zero—while the slope is positive and, at higher income levels, the band crosses zero, indicating no statistically detectable contemporaneous effect. Panel B conditions on lagged public debt and shows a downward slope: when debt is low, the effect is imprecisely estimated, but as debt ratios approach one-half of GDP and above, the entire confidence band falls below zero, implying unambiguously contractionary shocks in high-debt settings. Panel C conditions on trade openness and indicates that the effect is near zero for relatively closed economies but becomes significantly negative once openness exceeds roughly 70–75% of GDP, consistent with import leakage diluting the domestic demand impulse. Panel D conditions on the current-account balance and finds imprecision around deficits or balance, but a significantly negative effect once external surpluses emerge (about 2% of GDP and above), suggesting that surplus economies transmit unexpected environmental outlays more directly into output contractions within the year.

5.2 Subcomponents of Environmental Spending

To probe the aggregate results further, this subsection decomposes the innovation into its four COFOG categories—pollution abatement, environmental R&D, waste management, and wastewater—and re-estimates the reduced-form panel for each component in turn (country and year fixed effects; covariance matrix robust to heteroskedasticity, serial correlation, and cross-sectional dependence; shocks winsorised at the 1st–99th percentiles). The coefficients reported in Table 3 are interpretable as percentage-point changes in real GDP per-capita growth in response to a one–percentage-point-of-GDP surprise in the relevant spending category.

The estimates indicate that the short-run effects are not uniform across components. Shocks to pollution abatement are the most clearly contractionary: the point estimate is about –5.9 pp (significant at 10%). This accords with the compliance-driven, regulation-heavy nature of pollution control, which typically requires capital-intensive upgrades and can crowd out private activity on impact; evidence of near-term costs of environmental regulation is common in the micro literature (e.g., Greenstone, List and Syverson, 2012; Dechezleprêtre and Sato, 2017). For environmental R&D, the point estimate is even more negative (around –19.6 pp) but imprecise, reflecting the small budget share and lumpy, uncertain disbursement profile of R&D programmes—areas where payoffs tend to arrive with long lags (Hall, Mairesse and Mohnen, 2010; Bloom, Schankerman and Van Reenen, 2013). Shocks to waste management and wastewater are also negative (roughly –3.6 and –4.7 pp, respectively) but not statistically different from zero at conventional levels. These two categories are infrastructure-intensive and often import-heavy; planning, procurement, and execution lags mute the same-year demand impulse, and leakage abroad lowers multipliers in the short-run—patterns consistent with findings for public investment in developing economies (Ilzetzki, Mendoza and Végh, 2013; Bom and Ligthart, 2014).

Controls behave in line with the aggregate regressions. Lagged log income is negative and significant, pointing to conditional convergence. Trade openness is positive and economically meaningful (an additional 10 percentage points of openness is associated with about 0.4 pp faster growth), while inflation is negatively correlated with growth. Terms-of-trade improvements raise growth contemporaneously, and a higher current-account balance (toward surplus) is associated with weaker domestic absorption within the year. Debt and private credit are not robustly significant once these covariates and fixed effects are included.

Table 3. Reduced-form Panel Estimates: Growth Response to Environmental Spending Shocks by Component

Specification	(1)	(2)	(3)	(4)
Subcomponent	Pollution	R&D	Waste Mgmt	Waste Water
Environmental spending shock	-5.879*	-19.552	-3.610	-4.700
	(2.991)	(13.765)	(3.240)	(2.922)
Log real GDP per capita (t-1)	-4.967**	-4.911**	-4.901**	-4.900**
	(1.829)	(1.821)	(1.818)	(1.827)
Public debt (% of GDP) (t-1)	0.875	0.889	0.905	0.889
	(1.139)	(1.141)	(1.147)	(1.144)
Trade openness (% of GDP) (t-1)	0.041***	0.040***	0.040***	0.040***
	(0.013)	(0.013)	(0.013)	(0.013)
Consumer price inflation (annual %) (t-1)	-0.051**	-0.050**	-0.050**	-0.050**
	(0.019)	(0.019)	(0.019)	(0.019)
Private sector credit (% of GDP) (t-1)	-0.009	-0.008	-0.008	-0.009
	(0.009)	(0.009)	(0.009)	(0.009)
Change in terms of trade (t-1)	0.169***	0.169***	0.168***	0.170***
	(0.048)	(0.048)	(0.048)	(0.048)
Current-account balance (% of GDP) (t-1)	-0.073***	-0.073***	-0.071***	-0.073***
	(0.026)	(0.026)	(0.026)	(0.026)
Obs.	1035	1035	1035	1035

Note: The dependent variable is annual real GDP per-capita growth (percentage points). The key regressor is Environmental spending shock by subcomponent identified in the second row, a country-specific environmental-spending innovation constructed as the first-difference of residuals from country-by-country regressions of environmental protection outlays (share of GDP) on their own lag and lagged domestic fundamentals; the series is rescaled so that 1 unit = 1 percentage point of GDP and winsorised at the 1st–99th percentiles. Controls are lagged one year. All specifications include country fixed effects and year fixed effects; coefficients are therefore identified from within-country variation relative to the global co-movement in each year. Standard errors (in parentheses) are Spatial Correlation Consistent à la Driscoll–Kraay (1998). "Global," "AE," and "DEV" refer to the full sample, advanced economies, and EMDEs (developing/emerging), respectively; the panel is unbalanced, so coverage varies by country-year.
Asterisks denote significance: * p<0.10, ** p<0.05, *** p<0.01.

In combination with the aggregate results, the subcomponent evidence suggests that the short-run contraction is driven disproportionately by pollution-related expenditures, with R&D and waste-related categories pointing in the same direction but with weaker precision. This heterogeneity matters for policy design: categories that are compliance-oriented or capital-intensive are most likely to impose near-term output costs, even if they are precisely the items that generate medium-run productivity gains and resilience benefits once projects are completed (the logic behind the "Porter hypothesis"; Porter and van der Linde, 1995).

Because these mechanisms are inherently dynamic, the cross-sectional panel cannot speak to timing. The next section therefore turns to local projections, where we trace impulse responses for macroeconomic and environmental outcomes to shocks in total spending and in each subcomponent. This allows us to test whether the short-run costs documented here are offset by medium-run improvements in productivity, emissions intensity, and resilience, and whether the composition of spending shapes the trajectory of those responses.

5.3 Dynamic Local Projections

Figure 4 shows the unconditional response of real GDP per capita to a 1-percentage-point-of-GDP environmental spending shock in the global panel. The path is modestly hump-shaped: output falls on impact and reaches its trough around year 2 at about –0.05%, a magnitude that is economically small. The decline is statistically different from zero from impact through year 2; by year 3 the zero line enters the confidence band. From the mid-horizons onward the response gradually improves, crosses zero around year 5 and turns slightly positive at years 6–7, although these medium-run gains are not uniformly significant once uncertainty widens with horizon. Overall, the LP confirms the short-run contraction seen in the contemporaneous panel while leaving room for offsetting medium-run effects as projects come online—an issue we explore next with state-dependent and subcomponent LPs.

Figure 4. Unconditional Local Projection: Effect of a 1%-of-GDP Environmental Spending Shock on Real GDP per capita

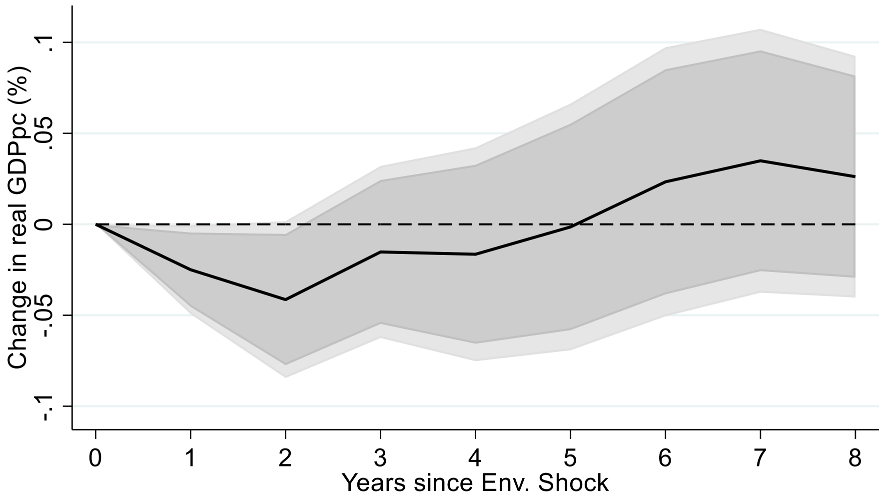

Note: The solid line plots the cumulative change in real GDP per capita (percent) at horizon h after a shock in year t (year 0 = shock year; year 1 = first year after the shock). Shocks are first-differenced residual innovations in government environmental protection outlays (COFOG), scaled so that one unit equals 1 percentage point of GDP and winsorised at the 1st–99th percentiles. Local projections include country fixed effects and year–horizon fixed effects, with AIC-selected dynamics: the shock enters contemporaneously; controls enter at t and t−1; and anticipation leads of the shock are included to address overlapping horizons. Shaded areas show 90% (dark) and 95% (light) confidence intervals computed with a covariance estimator robust to heteroskedasticity, serial correlation, and cross-sectional dependence (three-year window). Global, unbalanced panel.

Figure 5 splits the baseline local projections by income group. In advanced economies (left panel) the response to a 1-percentage-point-of-GDP environmental spending shock is negative on impact and remains below zero through mid-horizons, but the confidence bands overlap zero at all horizons. In other words, for AEs the short-run effect is negative but statistically indistinguishable from zero, and the medium-run recovery is likewise imprecise. In EMDEs (right panel) the pattern is sharper. Output falls in the first two years—the 90% band lies (mostly) below zero, making the early contraction borderline significant—and then gradually rebounds, crossing back toward zero around years 5–6, before easing again later in the window. This state of affairs mirrors the reduced-form and IV evidence: the global short-run contraction is largely driven by EMDEs. The interpretation is consistent with mechanisms more prevalent in developing economies—tighter financing and external constraints, higher import content of green capital goods, and execution lags—while AEs' deeper financial markets and project pipelines dampen near-term output costs.

Figure 5. Unconditional Local Projection by Income Group: Effect of a 1%-of-GDP Environmental Spending Shock on Real GDP per capita

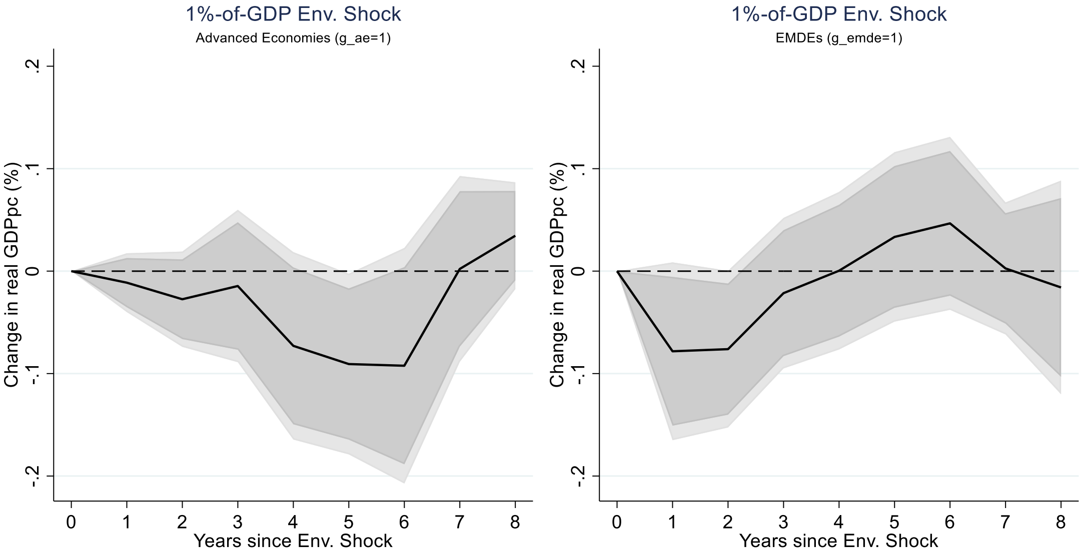

Note: The solid line plots the cumulative change in real GDP per capita (percent) at horizon h after a shock in year t (year 0 = shock year; year 1 = first year after the shock). Shocks are first-differenced residual innovations in government environmental protection outlays (COFOG), scaled so that one unit equals 1 percentage point of GDP and winsorised at the 1st–99th percentiles. Local projections include country fixed effects and year–horizon fixed effects, with AIC-selected dynamics: the shock enters contemporaneously; controls enter at t and t−1; and anticipation leads of the shock are included to address overlapping horizons. Shaded areas show 90% (dark) and 95% (light) confidence intervals computed with a covariance estimator robust to heteroskedasticity, serial correlation, and cross-sectional dependence (three-year window). Global, unbalanced panel.

Figure 6 decomposes the local projections by environmental subcomponent—pollution abatement, environmental R&D, waste management, and wastewater—using the same specification as in the baseline: the shock enters contemporaneously with anticipation leads; controls enter at ttt and t−1t-1t−1; and country and year–horizon fixed effects are included. The responses are cumulative changes in real GDP per capita to a 1-percentage-point-of-GDP surprise in each category, with 90% and 95% confidence bands. A pollution-abatement shock delivers the clearest near-term contraction. The response falls below zero on impact and remains statistically negative through roughly years 2–3, with the tightest bands among all subcomponents. This is the most precisely estimated component and lines up with the compliance- and capital-intensive nature of pollution control. An environmental R&D shock is also negative and statistically significant in the early years, before drifting back toward zero later in the horizon. Despite R&D's smaller budget share, the LP indicates a discernible short-run drag—consistent with front-loaded costs and deferred payoffs. For waste management, the response is negative and significant up to about year 3, after which it gradually reverts toward zero. The profile fits the idea that planning and procurement bottlenecks mute the same-year demand impulse. For wastewater, the initial contraction is somewhat more pronounced and longer-lived than for waste management; the band remains below zero into the mid-horizons before easing. This likely reflects large imported capital goods and execution lags typical of network investments.

Taken together, the subcomponent LPs confirm the aggregate pattern: short-run costs are present across categories but are best identified for pollution abatement, with R&D, waste management, and wastewater showing similarly negative early effects of smaller precision and shorter duration. In the next step, we examine whether these early costs are offset by medium-run gains and how the responses vary across country states, linking back to the heterogeneity documented above.

Figure 6. Unconditional Local Projection by Subcomponent: Effect of a 1%-of-GDP Environmental Spending Shock on Real GDP per capita

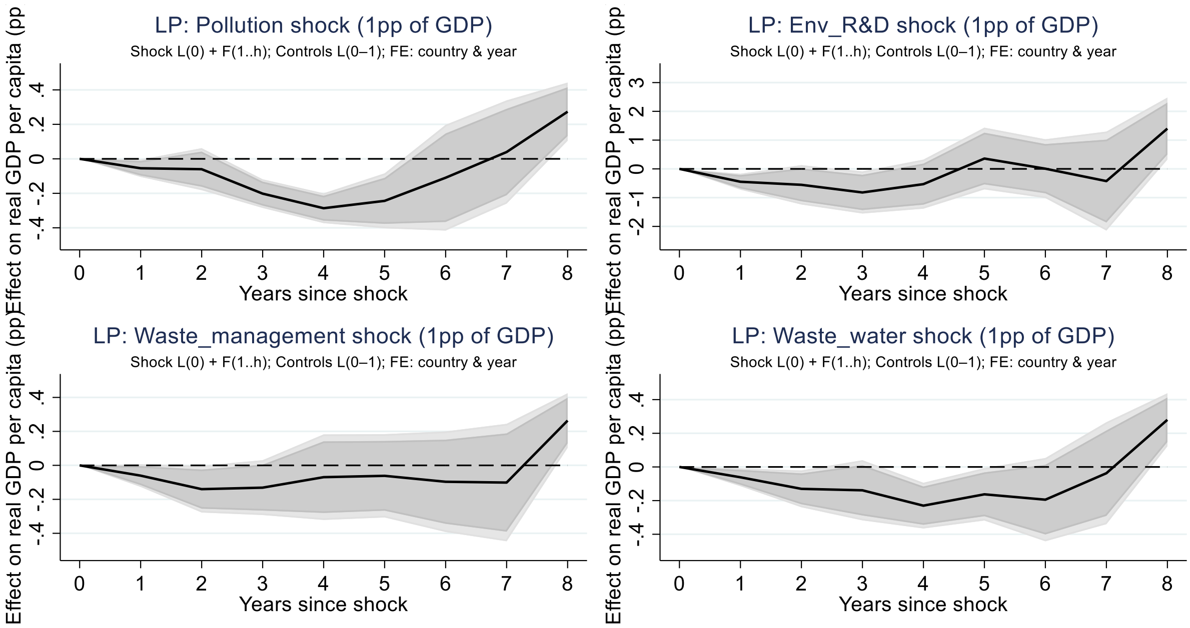

Note: The solid line plots the cumulative change in real GDP per capita (percent) at horizon h after a shock in year t (year 0 = shock year; year 1 = first year after the shock). Shocks are first-differenced residual innovations in government environmental protection outlays (COFOG), scaled so that one unit equals 1 percentage point of GDP and winsorised at the 1st–99th percentiles. Local projections include country fixed effects and year–horizon fixed effects, with AIC-selected dynamics: the shock enters contemporaneously; controls enter at t and t−1; and anticipation leads of the shock are included to address overlapping horizons. Shaded areas show 90% (dark) and 95% (light) confidence intervals computed with a covariance estimator robust to heteroskedasticity, serial correlation, and cross-sectional dependence (three-year window). Global, unbalanced panel.

In the AE subcomponent LPs (Figure A5 top panel), the clearest and most precisely estimated contractions follow pollution and wastewater shocks. Both profiles drop below zero on impact, reach their trough in the early-to-mid horizons, and their 90–95% confidence bands exclude zero for several years, indicating statistically significant short-run costs. By contrast, shocks to environmental R&D and waste management display similar shapes but their bands generally straddle zero throughout the window, so effects are not statistically different from zero. This pattern fits the idea that, in advanced economies with tighter environmental standards, compliance-intensive categories—pollution control and wastewater networks—front-load capital and adjustment costs, while R&D and routine waste operations are absorbed with less near-term macro impact. The EMDE subcomponent LPs (Figure A5 bottom panel) reveal a starker picture. The pollution shock again delivers the largest and most persistent contraction, with a pronounced early decline that remains statistically significant into the medium run before easing. Shocks to waste management and wastewater are also negative in the short run—their bands fall below zero in the first few years—followed by partial reversals later in the horizon. These dynamics are consistent with EMDE-specific frictions: higher import content of equipment, tighter financing and external constraints, and longer planning/procurement lags that depress near-term activity even when projects yield benefits later. Taken together, the income-group splits confirm that EMDEs drive the global short-run contraction, with the pollution component playing the leading role in both groups. This motivates the next set of exercises where we quantify how the timing and magnitude of these effects evolve across horizons and states—linking back to the aggregate LPs and to the heterogeneity by income, debt, openness, and external balance.

5.4 State Dependence in Local Projections

The preceding sections showed that aggregate and compositional shocks to environmental spending are followed by a small, short-run dip in activity and a gradual recovery. A natural next question is whether this average mask state dependence—i.e., whether the impact varies with prevailing macro conditions, financing constraints, external exposure, or institutional capacity. To investigate this, the local-projection framework is extended to allow the effect of the shock to depend on pre-determined country states, in the spirit of Auerbach and Gorodnichenko (2012) and Jordà and Taylor (2016).

We focus on three economically motivated conditioning variables. (i) Fiscal space, proxied by the public debt-to-GDP ratio, matters because high debt can tighten financing conditions, raise risk premia, and trigger offsetting consolidation—mechanisms that attenuate short-run multipliers. (ii) External exposure, measured by trade openness, shapes the import content of green investment and the scope for leakage of demand abroad. (iii) Adaptive capacity, captured by the ND-GAIN index (or, in separate specifications, its vulnerability/readiness components), summarises institutions, infrastructure and human capital that help convert spending into productivity and resilience gains. Each state variable is standardised to mean zero and unit variance to ensure comparability across specifications.

State dependence is implemented using a smooth-transition local projection (ST-LP). For each horizon h we estimate State-dependent LPs have been widely applied (Ramey and Zubairy, 2018; Alpanda et al., 2021; de Haan and Wiese, 2022). Plagborg-Møller and Wolf (2020) demonstrate that LPs and VARs yield the same IRFs in linear settings. However, LPs offer two advantages for our context: they easily accommodate state dependence in panels and do not require assumptions about regime duration or transition mechanisms. We estimate these nonlinear effects using a Smooth Transition Autoregressive (STAR) function within the local projections (LP) framework. As shown by Auerbach and Gorodnichenko (2012), this is equivalent to the STAR model of Granger and Teräsvirta (1993).[7] We estimate the following model:

$$\Delta y_{i,t+h} = \alpha_{i,h} + \delta_{t,h} + \beta_{0,h}^{L} F(z_{i,t}) * Shocki_{i,t}^{EP} + \beta_{0,h}^{H}(1 - F(z_{i,t})) * Shocki_{i,t}^{EP} + \sum_{c=0}^{1} \pi_{c,h} X_{i,t-c} + \sum_{m=0}^{h} \kappa_{m,h} Shocki_{i,t+m}^{EP} + u_{i,t+h} \tag{9}$$

where,

$$F(z_{it}) = \frac{\exp(-\gamma z_{it})}{1 + \exp(-\gamma z_{it})}, \quad \gamma > 0$$

with z_{it} being an indicator proxied by one of the variables mentioned above normalised to have zero mean and unit variance. The weights assigned to each regime vary between 0 and 1 according to the weighting function $F(.)$, so that $F(z_{it})$. The coefficients $\beta_{0,h}^{L}$ and $\beta_{0,h}^{H}$ are used to construct the IRFs and the associated confidence interval for low versus high states. They respectively capture the impact of environmental spending shocks at each horizon h in cases of low state ($F(z_{it})$≈1 when z goes to minus infinity) and high state (1-$F(z_{it})$)≈1 when z goes to plus infinity). The smoothness parameter (γ) governs the speed and steepness of transitions between regimes. Higher values of γ imply more abrupt, threshold-like shifts, while lower values generate smoother, gradual transitions. In our baseline specification, we calibrate γ = 1.5 in line with prior empirical work (see e.g., Auerbach and Gorodnichenko, 2012). To assess the robustness of our findings, we re-estimate the model using alternative values of γ=1.0 and γ=2.0. The resulting impulse responses remain stable across these values, confirming that our main results are not sensitive to this parameter choice. These robustness checks are available upon request.

The results in Figure 7 highlight a strong state-dependence in the effects of environmental spending shocks. Figure 7a conditions the response on public debt. When debt is low (left panel), the path remains close to zero at short horizons and turns positive in the medium-run, with point estimates rising above the unconditional

[7] This approach offers two main advantages. First, it directly tests whether fiscal shock effects differ across regimes, unlike models with business cycle interactions. Second, it allows shock responses to vary smoothly across a continuum of states, yielding more stable and precise impulse responses than regime-specific structural VARs.

band from about year 4 onward. This is consistent with the idea that ample fiscal space relaxes financing constraints and keeps risk premia contained, allowing green projects to proceed without offsetting consolidation or crowding-out; the medium-run gains then materialise as assets come on line. By contrast, when debt is high (right panel) the response is significantly negative on impact and through the early years, indicating that additional outlays are met with tighter borrowing conditions, precautionary private saving, or expectations of future fiscal adjustment. These patterns mirror state-dependent fiscal multipliers reported in the literature, where effects are larger—and less contractionary—when financing conditions are slack and credibility is high (Auerbach and Gorodnichenko, 2012; Corsetti, Meier and Müller, 2012).

Figure 7a. Conditional Local Projections by Fiscal Space (Public Debt): Effect of a 1%-of-GDP Environmental Spending Shock on Real GDP per capita

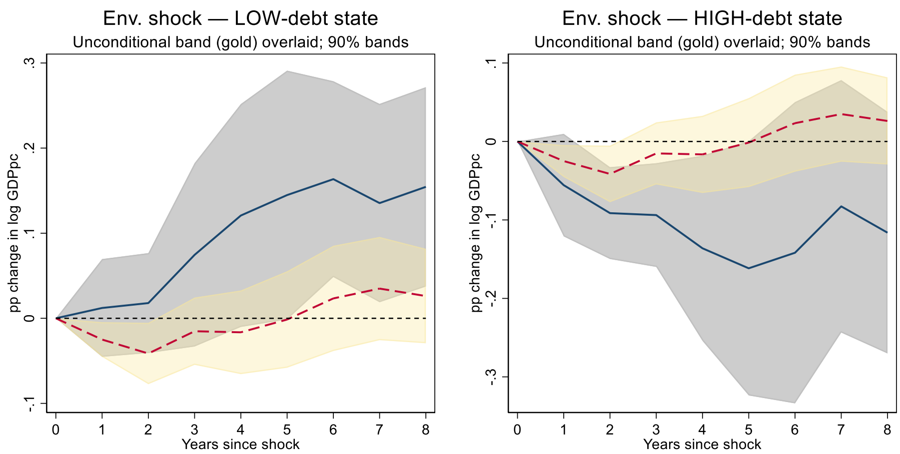

Note: Shocks are first-differenced residual innovations in government environmental-protection outlays (COFOG), scaled so that one unit equals 1 percentage point of GDP and winsorised at the 1st–99th percentiles. Year 0 is the shock year; Year 1 is the first year after the shock. The solid blue line plots the conditional impulse response for the state indicated above (e.g., low vs. high debt, low vs. high trade openness, low vulnerability vs. high readiness). The dashed red line shows the unconditional baseline response from the global specification. Dark-gray bands are 90% confidence intervals for the conditional response; gold bands are 90% confidence intervals for the unconditional response. All responses come from local projections with country and year–horizon fixed effects, the shock entering contemporaneously with anticipation leads, and controls at t and t–1; standard errors are robust to heteroskedasticity, serial correlation, and cross-sectional dependence (three-year window).

Figure 7b conditions on trade openness. In low-openness states (left panel) a negative and statistically significant short-run effect is observed, with a subsequent rebound that converges toward the unconditional path at medium horizons. One interpretation is that in relatively closed economies, the initial compliance and execution costs dominate and there is limited offset from export demand or supply-chain spillovers. In very open economies (right panel) the response is close to zero throughout, suggesting that import leakage and flexible external adjustment blunt the domestic demand impulse but also cushion the downturn—

yielding small net effects overall. This attenuation of multipliers with openness echoes classic open-economy findings (Ilzetzki, Mendoza and Végh, 2013; Corsetti et al., 2012).

Figure 7b. Conditional Local Projections by External Exposure (Trade Openness): Effect of a 1%-of-GDP Environmental Spending Shock on Real GDP per capita

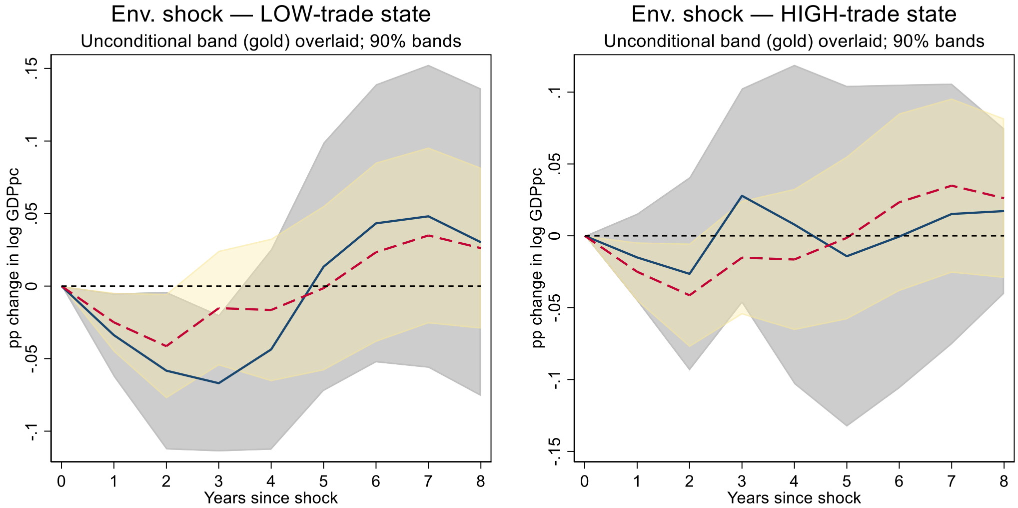

Note: Shocks are first-differenced residual innovations in government environmental-protection outlays (COFOG), scaled so that one unit equals 1 percentage point of GDP and winsorised at the 1st–99th percentiles. Year 0 is the shock year; Year 1 is the first year after the shock. The solid blue line plots the conditional impulse response for the state indicated above (e.g., low vs. high debt, low vs. high trade openness, low vulnerability vs. high readiness). The dashed red line shows the unconditional baseline response from the global specification. Dark-gray bands are 90% confidence intervals for the conditional response; gold bands are 90% confidence intervals for the unconditional response. All responses come from local projections with country and year–horizon fixed effects, the shock entering contemporaneously with anticipation leads, and controls at t and t–1; standard errors are robust to heteroskedasticity, serial correlation, and cross-sectional dependence (three-year window).

Figure 7c conditions on adaptive capacity using the ND-GAIN family of indicators. ND-GAIN has two complementary pillars: Vulnerability (exposure and sensitivity across food, water, health, ecosystems, human habitat, and infrastructure) and Readiness, which we label “resilience” in the figure for readability. Readiness embeds an explicitly institutional block—governance readiness (regulatory quality, rule of law, control of corruption)—alongside economic and social readiness; higher values therefore proxy stronger institutions and implementation capacity in exactly the sense relevant for executing green projects (planning, procurement, permitting, contract enforcement). In the low-vulnerability state (left panel), a negative and statistically significant short-run effect is followed by a steady recovery that eventually rises above the unconditional band. This pattern is consistent with countries that face fewer service and infrastructure gaps absorbing the initial compliance and capital costs more efficiently—and then converting projects into productivity and resilience gains once assets come online. In the high-resilience (high-readiness) state (right panel), the short-run contraction is again significant, but the turnaround is faster and the medium-run profile is more favourable than in the unconditional case. Interpreted through the ND-GAIN lens, stronger governance and market institutions shorten execution lags, reduce cost overruns, and ease financing constraints, so the mechanical initial dip (front-loaded compliance and lumpy investment) dissipates sooner and medium-term payoffs are larger. This timing is in line with evidence on delayed returns to R&D- and capital-intensive public programmes (Hall, Mairesse and Mohnen, 2010; Bloom, Schankerman and Van Reenen, 2013).

Across debt, openness, and adaptive-capacity states, a coherent picture emerges: short-run costs are pervasive but amplified when fiscal space is tight, non-salient in very open economies, and more quickly reversed where institutions are stronger. These results align with state-dependent multiplier theory and help explain why the global average masks substantial heterogeneity across country circumstances.

Figure 7c. Conditional Local Projections by Adaptive Capacity (ND-GAIN: Low Vulnerability vs. High Readiness): Effect of a 1%-of-GDP Environmental Spending Shock on Real GDP per capita

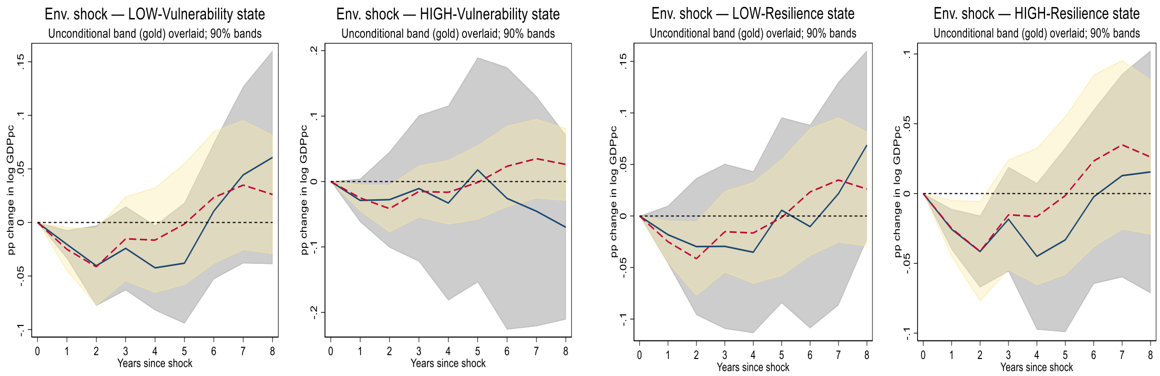

Note: Shocks are first-differenced residual innovations in government environmental-protection outlays (COFOG), scaled so that one unit equals 1 percentage point of GDP and winsorised at the 1st–99th percentiles. Year 0 is the shock year; Year 1 is the first year after the shock. The solid blue line plots the conditional impulse response for the state indicated above (e.g., low vs. high debt, low vs. high trade openness, low vulnerability vs. high readiness). The dashed red line shows the unconditional baseline response from the global specification. Dark-gray bands are 90% confidence intervals for the conditional response; gold bands are 90% confidence intervals for the unconditional response. All responses come from local projections with country and year–horizon fixed effects, the shock entering contemporaneously with anticipation leads, and controls at t and t–1; standard errors are robust to heteroskedasticity, serial correlation, and cross-sectional dependence (three-year window).

5.5 Robustness and Sensitivity

Appendix Figure A6 assesses how the local-projection point estimates vary with alternative control sets. Each thin line corresponds to one specification using a different subset of lagged controls (entered at t and t−1); the three highlighted paths summarise the cases with no controls, income only (lagged log GDP per capita), and the full set. All specifications include country and year–horizon fixed effects, use the same shock timing (contemporaneous shock plus anticipation leads), and are estimated with a covariance matrix robust to heteroskedasticity, serial correlation, and cross-sectional dependence. A common pattern emerges across specifications. The response is negative on impact, bottoms out around year 2, and then rebounds, crossing zero roughly in the middle of the window and turning mildly positive toward years 6–8. Differences across control sets are second order: adding the full controls nudges the path only slightly (a marginally shallower trough and a similar recovery), while the income-only path lies between the “no controls” and “full controls” lines. No specification reverses the sign of the near-term contraction or the subsequent improvement. The dispersion of lines is modest relative to the uncertainty bands reported in the baseline figure, indicating that the dynamic profile is not an artifact of control choice but is instead driven by within-country variation once common shocks are removed. Overall, this sensitivity check supports the baseline identification: the short-run dip and gradual recovery are stable to reasonable perturbations of the macro control set, reinforcing the interpretation of the local-projection results presented in Section 5.3.

To further assess whether the residual-based environmental spending shocks truly behave as unanticipated innovations rather than systematic responses to macro conditions, we run a suite of predictability and timing checks that mirror our empirical setup (country and year fixed effects; covariance matrix robust to heteroskedasticity, serial correlation, and cross-sectional dependence).

Panel Granger non-causality. Using the Dumitrescu–Hurlin (2012) panel test over 1–3 lags, we ask whether past macro variables help forecast the continuous shock series. Across growth, inflation, public debt, trade openness, private credit, the current-account balance, and the terms of trade, the null of no Granger causality is not rejected. Results are similar when the test is run on a balanced subpanel and when emphasis is placed on variables not used in the first-stage residualisation (so as not to test, by construction, what has already been purged). This supports the view that shocks are not systematically preceded by macro developments.

Predictability regressions. We regress the shock on three lags of the same macro covariates in a fixed-effects panel with year dummies and SCC errors. Individual coefficients are small and statistically indistinguishable from zero, and a joint F-test of all lagged predictors fails to reject. In short, the shocks are hard to forecast using available macro information.

Pre-trends around large shocks. To visualise timing, we construct an event-study around “large” innovations—years in the top decile of the absolute shock (one event per country, earliest occurrence)—and estimate local projections of real GDP per-capita growth on lead/lag dummies controlling for lagged outcomes and macro covariates. Pre-event coefficients (–3, –2, –1) are jointly indistinguishable from zero, while movements appear at the event and just after. This pattern indicates flat pre-trends and suggests that shock timing is not chosen in anticipation of near-term macro changes.

Placebo timing. Regressions that replace the contemporaneous innovation with future shocks (e.g., leads at +2 and +3) produce no significant effects on current growth, reinforcing the interpretation that the dynamic responses we estimate are not driven by foresight.

Orthogonality to contemporaneous levels. A regression of the shock on the level of environmental spending (with country fixed effects and year dummies) yields a near-zero, insignificant association, consistent with the construction that removes persistence and low-frequency drift and with the visual evidence in the appendix scatter. Hence, the measured innovations are not simply scaled versions of the spending level.

Taken together, these diagnostics point in the same direction: the residual-based environmental spending shocks are orthogonal to recent and foreseeable macro conditions and are not timed to future outcomes. This strengthens the credibility of our identification strategy and the effects and discussions reported earlier.

6. Conclusions and Policy Implications

This paper set out to measure how unanticipated innovations in environmental protection outlays propagate through the macroeconomy and how those effects vary across countries and spending categories. We constructed country-specific, residual-based shocks to COFOG-05 expenditure (scaled so 1 unit = 1 pp of GDP) by purging persistence and lagged domestic fundamentals, and then brought those innovations to reduced-form panels, instrumental-variable specifications, and local projections—including state-dependent LPs that condition on fiscal space, external exposure, and adaptive capacity (ND-GAIN).

The evidence is coherent and lines up closely with the model's predictions. In the model, green outlays entail front-loaded compliance and investment costs that can crowd out private activity or bind financing constraints on impact, while productivity and resilience benefits arrive as assets come online and firms/households adjust. Empirically, the global reduced form shows that a 1-pp-of-GDP innovation is associated with lower contemporaneous per-capita growth (≈ –2.6 pp), with larger IV magnitudes (≈ –4 pp globally; ≈ –7 pp in EMDEs). The baseline LP traces a small, short-run dip that troughs around year 2 (~ –0.05%), is statistically different from zero through year 2, and then fades and slightly turns positive by years 6–7, consistent with the model's medium-run channel. Composition matters in the manner anticipated: pollution-abatement shocks deliver the clearest short-run contraction; wastewater and waste management are also negative on impact but less precisely estimated; environmental R&D shows an early drag with deferred payoffs—exactly the ordering implied by compliance intensity, import content, and gestation lags.

Heterogeneity is large and economically meaningful. EMDEs drive the global short-run contraction (AEs are negative but statistically indistinct from zero), and state dependence operates as the model would suggest: high debt amplifies the early decline (tighter financing/credible consolidation expectations), very open economies display attenuated net effects (import leakage and external adjustment cushion demand), and higher ND-GAIN readiness—which embeds an explicit governance block and thus proxies implementation capacity—does not eliminate the initial dip but shortens its duration and raises medium-run payoffs. Robustness checks (panel Granger non-causality, predictability regressions, pre-trend event studies, placebos, and orthogonality to spending levels) support the shocks' exogeneity and timing.

These findings carry clear policy implications. First, sequencing matters: pairing environmental outlays with steps that expand fiscal space (e.g., concessional financing, credible medium-term anchors) reduces near-term crowding-out. Second, composition matters: compliance-intensive categories (pollution control, wastewater networks) should be embedded in procurement and delivery models that minimise execution lags and cost overruns; for infrastructure with high import content, local-content, skills, and supplier-development policies can limit leakage. Third, state-contingent implementation is advisable: where debt is high or institutions are weaker, emphasise readiness-enhancing reforms (permits, procurement, O&M) and phased roll-outs to smooth the near-term macro footprint. Finally, the medium-run upside—in productivity, energy efficiency, and resilience—argues for patient evaluation windows and performance-based budgeting rather than judging programmes on impact-year GDP alone.

Two limitations merit mention without undermining the core results. First, while residualisation plus differencing yields plausibly exogenous innovations, measurement error in spending and coarse COFOG subcomponents can attenuate precision, especially for R&D. Second, LPs recover average responses under parsimonious dynamics; very long-horizon benefits (beyond eight years) and general-equilibrium reallocation are not fully captured. That said, the results are stable across control sets, inference procedures, and foresight tests, and they map tightly to theory.

Promising avenues for future work include: linking project-level or procurement microdata to macro outcomes to unpack execution frictions; studying private crowd-in (green FDI, firm investment, credit) and supply-chain adjustments; integrating carbon-pricing and regulatory reforms to quantify policy packages; exploiting exogenous funding shocks (e.g., grants, climate finance) or narrative events for sharper identification; and extending the horizon and outcome set (air quality, adaptation damages avoided) to better capture long-run returns. Together, these steps would deepen our understanding of when and how green public spending delivers both near-term macro manageability and durable environmental gains.

7. Implications for Central Banks

The findings also have implications for central banks. Because green public spending can generate short-run output costs that vary substantially with fiscal space and implementation capacity, monetary authorities should incorporate the scale, composition, and timing of green fiscal measures into their assessment of the policy mix rather than treating them as purely structural fiscal interventions. Where fiscal space is limited, abrupt or poorly sequenced green investment may amplify macroeconomic and financial vulnerabilities through weaker activity, tighter financing conditions, and pressures on sovereign and private sector balance sheets, warranting close monitoring from a financial stability perspective. At the same time, because the adverse effects are largely front-loaded while productivity and resilience gains emerge gradually, monetary policy should avoid responding mechanically to temporary adjustment effects and instead distinguish them from more persistent changes in inflation, demand, and potential output.

More broadly, the green transition could also reshape monetary policy transmission in ways that conventional macroeconomic frameworks do not fully capture. By altering production structures, investment patterns, and energy use, environmental spending can affect demand, supply conditions, inflation expectations, and financial conditions simultaneously. Central banks should therefore account for these evolving interactions when assessing the effectiveness and transmission of monetary policy.

Overall, these considerations suggest that the green transition should become an increasingly integrated component of surveillance and analysis by central banks. Greater attention to transition-related investment, financing conditions, and sectoral adjustment can improve the identification of underlying inflationary and growth pressures and reduce the risk of policy miscalibration. Over time, incorporating climate-related fiscal developments into forecasting and risk assessment frameworks will become more important as environmental policies influence capital allocation, productivity, and financial stability. This can help central banks to maintain price and financial stability while supporting an orderly transition to a lower-carbon economy.

References

Acemoglu, D., Aghion, P., Bursztyn, L., & Hémous, D. (2012). The environment and directed technical change. *American Economic Review*, *102*(1), 131–166. doi.org

Alpanda, S., Fernández, A., Jara-Figueroa, J. M., & Giuliodori, M. (2021). The macroeconomic effects of fiscal consolidations with and without monetary accommodation. *Journal of Monetary Economics*, *118*, 1–20. doi.org

Auerbach, A. J., & Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, *4*(2), 1–27. doi.org

Auerbach, A. J., & Gorodnichenko, Y. (2013). Fiscal multipliers in recession and expansion. In A. Alesina & F. Giavazzi (Eds.), *Fiscal policy after the financial crisis* (pp. 63–98). University of Chicago Press.

Auerbach, A. J., & Gorodnichenko, Y. (2017). Fiscal stimulus and fiscal sustainability. In R. Hemming, M. Kell, & A. Schimmelpfennig (Eds.), *Navigating fiscal challenges in the wake of the global crisis*. International Monetary Fund.

Barro, R. J. (1990). Government spending in a simple model of endogenous growth. *Journal of Political Economy*, *98*(5), S103–S125. doi.org

Batini, N., Plotnikov, D., & Gupta, S. (2021). Fiscal policies for climate adaptation: Macroeconomic effects and distributional implications. *IMF Economic Review*, *69*(4), 1–48. doi.org

Baxter, M., & King, R. G. (1993). Fiscal policy in general equilibrium. *American Economic Review*, *83*(3), 315–334.

Bom, P. R. D., & Ligthart, J. E. (2014). What have we learned from three decades of research on the productivity of public capital? *Journal of Economic Surveys*, *28*(5), 889–916. doi.org

de Haan, J., & Wiese, R. (2022). The state-dependent effects of fiscal policy: A review. *Journal of Economic Surveys*, *36*(3), 685–719. doi.org

Granger, C. W. J., & Teräsvirta, T. (1993). *Modelling nonlinear economic relationships*. Oxford University Press.

Hallegatte, S., Vogt-Schilb, A., Bangalore, M., & Rozenberg, J. (2016). *Unbreakable: Building the resilience of the poor in the face of natural disasters*. World Bank.

Hepburn, C., O'Callaghan, B., Stern, N., Stiglitz, J., & Zenghelis, D. (2020). Will COVID-19 fiscal recovery packages accelerate or retard progress on climate change? *Oxford Review of Economic Policy*, *36*(S1), S359–S381. doi.org

Heutel, G. (2012). How should environmental policy respond to business cycles? Optimal policy under persistent productivity shocks. *Review of Economic Dynamics*, *15*(2), 244–264. doi.org

Ilzetzki, E., Mendoza, E. G., & Végh, C. A. (2013). How big (small?) are fiscal multipliers? *Journal of Monetary Economics*, *60*(2), 239–254. doi.org

International Monetary Fund. (2022). *Fiscal policy for a resilient and green recovery*.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, *95*(1), 161–182. doi.org

Jordà, Ò., & Taylor, A. M. (2016). The time for austerity: Estimating the average treatment effect of fiscal policy. *Economic Journal*, *126*(590), 219–255. doi.org

Kraay, A. (2014). Government spending multipliers in developing countries: Evidence from lending by official creditors. *American Economic Journal: Macroeconomics*, *6*(4), 170–208. doi.org

Mertens, K., & Ravn, M. O. (2012). Empirical evidence on the aggregate effects of anticipated and unanticipated US tax policy shocks. *American Economic Journal: Economic Policy*, *4*(2), 145–181. doi.org

Nordhaus, W. D. (2014). *The climate casino: Risk, uncertainty, and economics for a warming world*. Yale University Press.

Pereira, A. M., & Pereira, R. M. (2017). On the effects of public investment on economic performance. *Journal of Economic Development*, *42*(2), 1–25.

Plagborg-Møller, M., & Christian, K. W. (2020). Local projections and VARs estimate the same impulse responses. *Econometrica*, *88*(2), 955–980. doi.org

Ramey, V. A. (2011). Identifying government spending shocks: It's all in the timing. *Quarterly Journal of Economics*, *126*(1), 1–50. doi.org

Ramey, V. A. (2020). Ten years after the financial crisis: What have we learned from the renaissance in fiscal research? *Journal of Economic Perspectives*, *33*(2), 89–114. doi.org

Ramey, V. A., & Zubairy, S. (2018). Government spending multipliers in good times and in bad: Evidence from U.S. historical data. *Journal of Political Economy*, *126*(2), 850–901. doi.org

Romer, C. D., & Romer, D. H. (2010). The macroeconomic effects of tax changes: Estimates based on a new measure of fiscal shocks. *American Economic Review*, *100*(3), 763–801. doi.org

Teulings, C. N., & Zubanov, N. (2014). Is economic recovery a myth? Robust estimation of impulse responses. *Journal of Applied Econometrics*, *29*(3), 497–514. doi.org

APPENDIX

Table A1. Country List and Income Groups

Advanced Economies (AEs)
United States, Austria, Belgium, Denmark, Luxembourg, Norway, Sweden, Switzerland, Canada, Japan, Finland, Greece, Iceland, Ireland, Malta, Spain, Cyprus, Israel, Hong Kong SAR, Korea (Republic of), Singapore, Czech Republic, Slovak Republic, Estonia, Slovenia.

Emerging Markets & Developing Economies (EMDEs)
South Africa, Argentina, Bolivia, Brazil, Chile, Colombia, Dominican Republic, Ecuador, El Salvador, Guatemala, Haiti, Mexico, Nicaragua, Paraguay, Peru, Uruguay, Venezuela (Republica Bolivariana de), Bahamas (The), Belize, Bahrain, Iran (Islamic Republic of), Iraq, Kuwait, Lebanon, Oman, Saudi Arabia, Syria, Bangladesh, Bhutan, Brunei Darussalam, Sri Lanka, India, Timor Leste, Malaysia, Maldives, Nepal, Pakistan, Thailand, Vietnam, Djibouti, Algeria, Burundi, Cameroon, Cape Verde, Chad, Comoros, Congo (Republic of), Benin, Eritrea, Gabon, Gambia (The), Ghana, Guinea, Kenya, Lesotho, Libya, Madagascar, Mali, Mauritania, Morocco, Mozambique, Niger, Zimbabwe, Senegal, Sierra Leone, Namibia, Sudan, Swaziland, Togo, Tunisia, Uganda, Burkina Faso, Solomon Islands, Vanuatu, Samoa, Tonga, Albania, Kazakhstan, Kyrgyz Republic, Russian Federation, Ukraine, Montenegro, Hungary, Mongolia, Croatia, Macedonia (FYR), Bosnia & Herzegovina, Poland, Romania.

Figure A1. Distribution of Environmental Spending Shocks

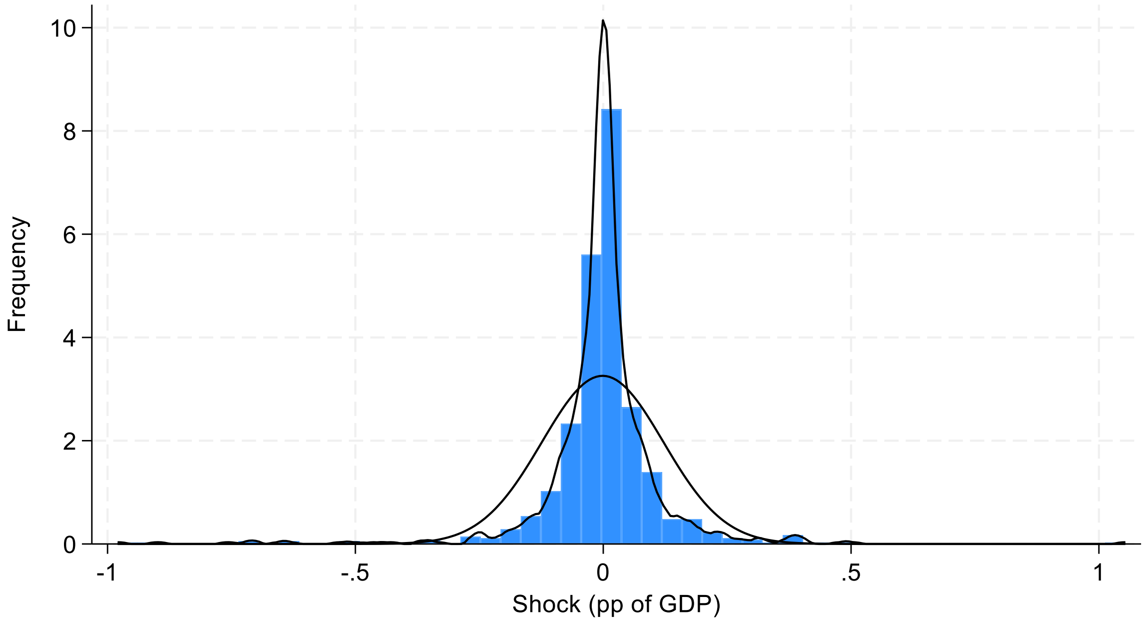

Note: The figure displays the empirical distribution of environmental protection spending shocks (percentage points of GDP). The histogram is overlaid with a kernel density estimate. Shocks are constructed as the first-differenced residuals from the country-specific orthogonalisation of environmental spending described in Section 4.1 and are expressed in percentage points of GDP.

Figure A2. Distribution of Environmental Spending Shocks by Decade

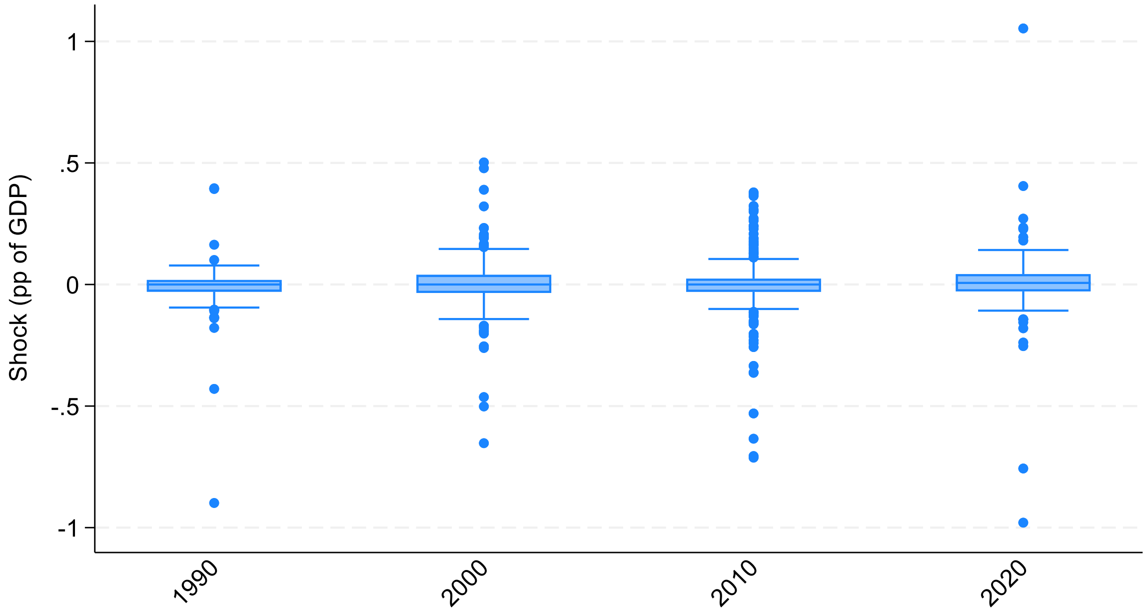

Note: The figure shows the distribution of environmental protection spending shocks (percentage points of GDP) by decade using box–dot plots. Each dot represents a country-year observation. The boxes display the interquartile range, with whiskers extending to the most extreme observations within 1.5 times the interquartile range. Shocks are constructed as described in Section 4.1.

Figure A3. Time Series of Environmental Spending Shocks: United Kingdom

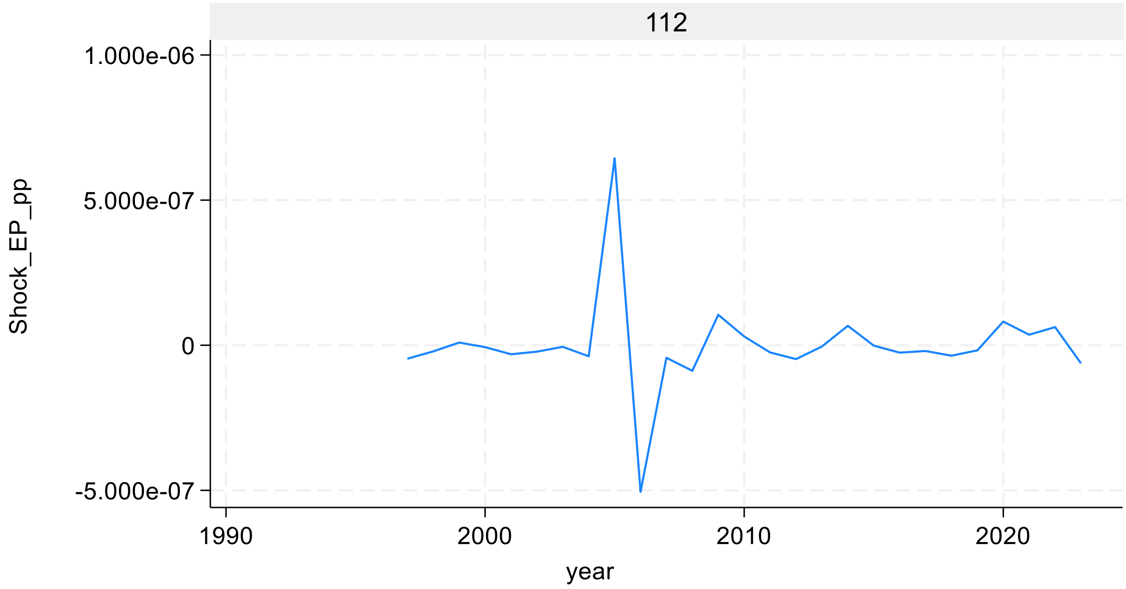

Note: The figure reports the time series of the environmental protection spending shock for the United Kingdom. The shock is measured in percentage points of GDP and is constructed as the first-differenced residual from the country-specific orthogonalisation of environmental spending, as described in Section 4.1. The horizontal axis denotes calendar years.

Figure A4. Contemporaneous Relationship Between Environmental Spending and the Shock

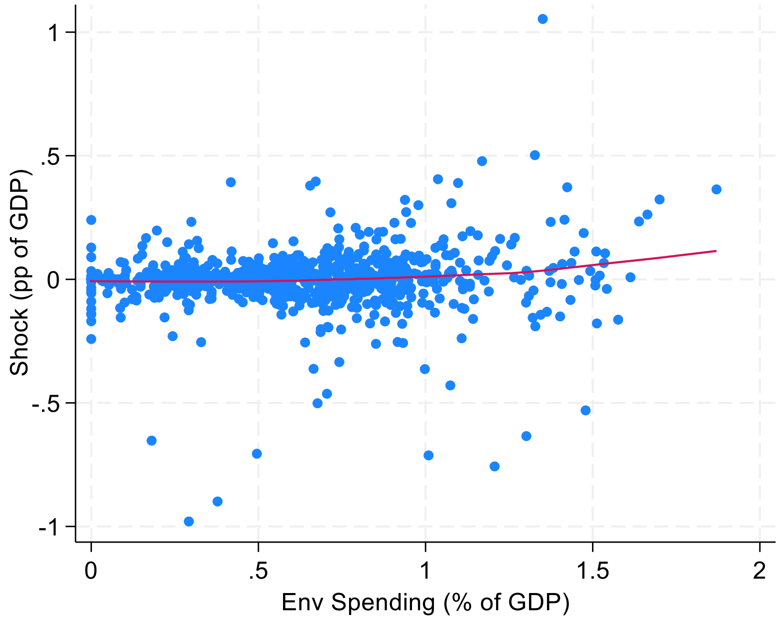

Note: The figure plots the contemporaneous relationship between the environmental spending shock (percentage points of GDP) and the level of environmental protection expenditure (percent of GDP). Each point represents a country-year observation. The solid line corresponds to a nonparametric lowess fit. The shock is constructed as the first-differenced residual from the country-specific orthogonalisation described in Section 4.1 and is scaled so that one unit equals one percentage point of GDP.

Figure A5. Unconditional Local Projection by Subcomponent and Income Group: Effect of a 1%-of-GDP Environmental Spending Shock on real GDP per capita

Panel A. AE

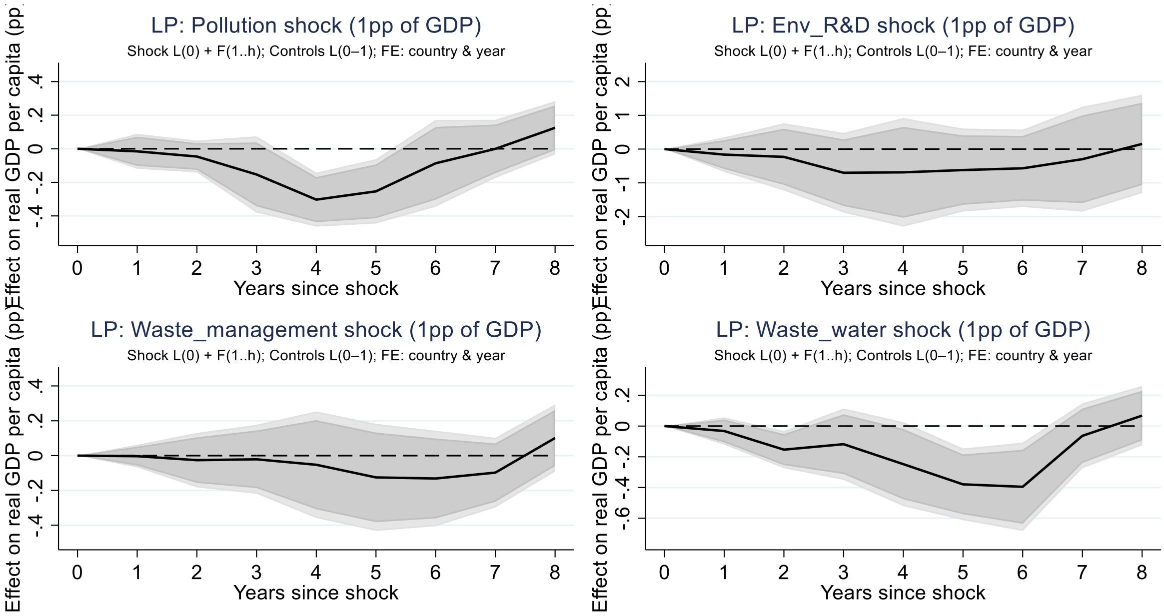

Panel B. EMDE

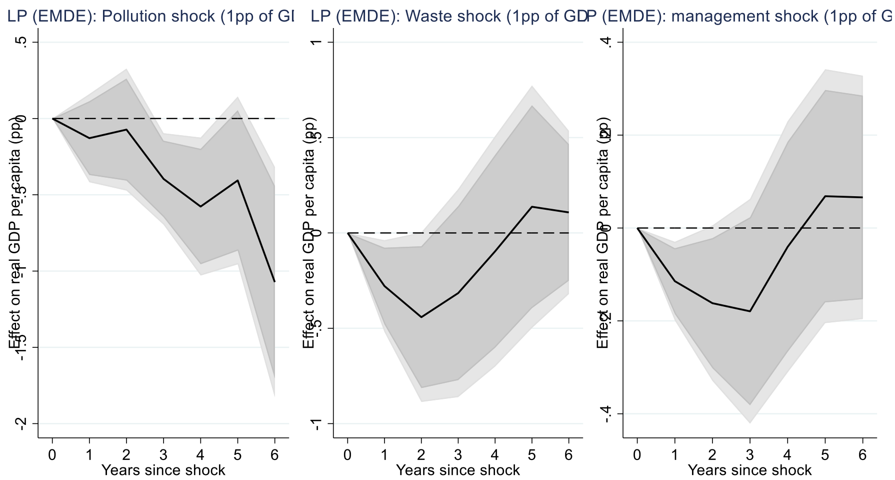

Note: The solid line plots the cumulative change in real GDP per capita (percent) at horizon h after a shock in year t (year 0 = shock year; year 1 = first year after the shock). Shocks are first-differenced residual innovations in government environmental protection outlays (COFOG), scaled so that one unit equals 1 percentage point of GDP and winsorised at the 1st–99th percentiles. Local projections include country fixed effects and year–horizon fixed effects, with AIC-selected dynamics: the shock enters contemporaneously; controls enter at t and t−1; and anticipation leads of the shock are included to address overlapping horizons. Shaded areas show 90% (dark) and 95% (light) confidence intervals computed with a covariance estimator robust to heteroskedasticity, serial correlation, and cross-sectional dependence (three-year window). Global, unbalanced panel.

Figure A6. Sensitivity of Local Projections Point Estimates to Control Sets

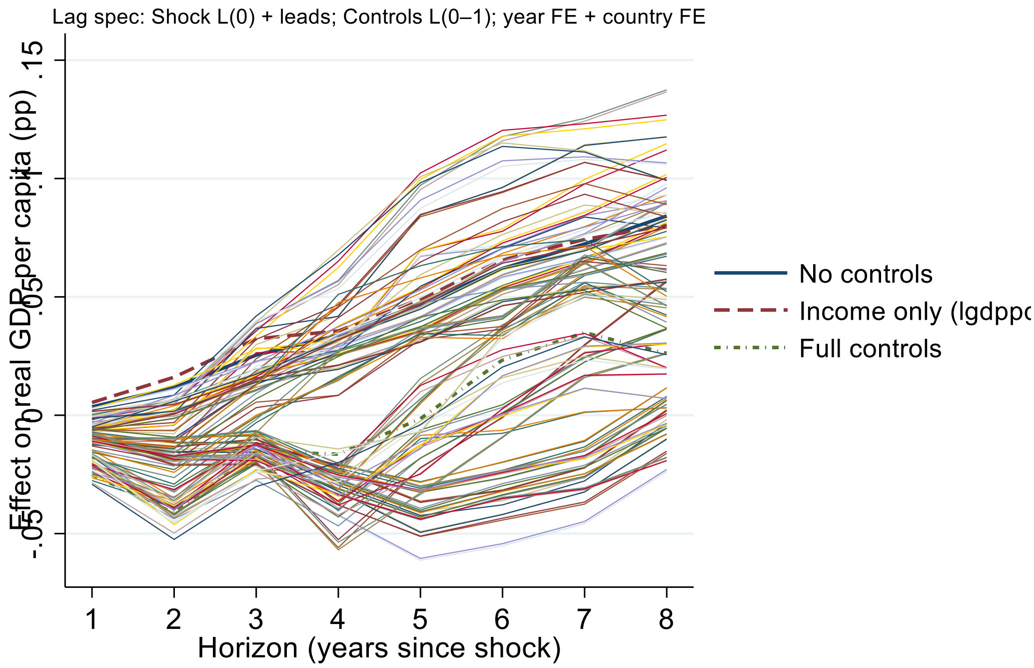

Note: The figure plots cumulative impulse responses of real GDP per capita to a 1-percentage-point-of-GDP innovation in environmental protection spending under alternative control specifications. Each thin line corresponds to one specification using a different subset of lagged macroeconomic controls. The three highlighted paths represent the cases with no controls, income only (lagged log GDP per capita), and the full control set. Local projections include country and year fixed effects. The shock enters contemporaneously with anticipation leads, and controls enter at t and t–1. The horizontal axis reports the horizon in years since the shock.