NBER WORKING PAPER SERIES

THE RETURN TO ADAPTATION IN A CHANGING CLIMATE

Harrison Hong
Serena Ng
Jiangmin Xu

Working Paper 33824
http://www.nber.org/papers/w33824

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2025, revised May 2026

We thank Benjamin Olken, Robin Burgess, Maryam Farboodi, Allan Hsiao, Marshall Burke, Adrien Bilal, Mark Watson, Chad Jones, Nobuhiro Kiyotaki, and seminar participants at the London School of Economics Environment Week, Stanford GSB Climate Conference, Columbia University, Econometric Society Meetings (Ho Chi Minh), Insper, Australasian Banking and Finance Conference, Macquarie University, University of Sydney, and Sabanci University (Istanbul) for helpful comments. Ng acknowledges financial support from the National Science Foundation (SES 2018369). Hong acknowledges financial support from a Guggenheim Fellowship. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Return to Adaptation in a Changing Climate
Harrison Hong, Serena Ng, and Jiangmin Xu
NBER Working Paper No. 33824
May 2025, revised May 2026
JEL No. O1, O40, O47, Q50, Q54, Q56

ABSTRACT

We quantify the economic return to adaptation in a changing climate. Our approach exploits the fact that extreme-weather arrivals are nonhomogeneous: we first estimate each country's time-varying disaster arrival rate, then test whether per-disaster damage declines when risk is elevated—a key feature of risk-dependent adaptation. We apply this to tropical cyclones across 48 countries over 1980–2019. Indeed, cyclones do significantly less damage when the arrival rate has risen above its historical baseline. The finding survives placebo permutations, lagged-disaster controls, and two-way-fixed-effects critiques. Welfare gains reach several percent of GDP for the most exposed economies and scale sharply under climate projections.

Harrison Hong
Columbia University
Department of Economics
and NBER
hh2679@columbia.edu

Serena Ng
Columbia University
Department of Economics
and NBER
Serena.Ng@columbia.edu

Jiangmin Xu
Peking University
Department of Finance
jiangminxu@gsm.pku.edu.cn

1 Introduction

The economic impact of extreme weather is large and rising. A growing empirical literature confirms that cyclones, heatwaves, and other climate-linked shocks meaningfully damage long-run economic growth (see Dell, Jones, and Olken (2014) for a review).[1] As climate change pushes the frequency and intensity of these events upward (National Academies of Sciences, Engineering, and Medicine, 2016), the policy-relevant question is how much of the projected damage countries can be offset by adaptation—by adjusting economic activity as the underlying disaster risk evolves. The dynamic return to adaptation is a key input for climate-economy modeling, fiscal planning under climate stress, and the pricing of catastrophe risk, yet it remains imprecisely quantified (Bouwer, Crompton, Faust, Höppe, and Pielke Jr, 2007).

This paper proposes an empirical framework that quantifies the return to adaptation as a country's disaster risk evolves over time. The economic logic is simple. If a country's risk of being struck by a disaster has been drifting upward, the optimizing agents have a sharper incentive to adapt than they did when the risk was lower. The damage caused by any individual disaster should therefore be smaller when current risk is high than when it is low. The return to risk-dependent adaptation is the rate at which the per-disaster damage falls as the country's risk rises.

To make this idea operational we adopt a two-stage estimation strategy. In the first stage, we model each country's disaster arrival process as a nonhomogeneous Poisson process (NHPP) with a country-specific time-varying intensity λ_{it}, and estimate the parameters of the intensity path by maximum likelihood. We allow the intensity to follow a constant, linear-trend, or quadratic-trend specification and test among them. In the second stage, we relate the country-year growth rate to the disaster indicator D_{it}, the estimated risk path $\hat{\lambda}_{i,t-1}$, and their interaction in a panel regression. The interaction coefficient is the empirical analog of the slope of an adaptation function, in the spirit of the continuous-time growth model with adaptation in Hong, Wang, and Yang (2023): a positive interaction means the per-disaster damage is smaller in country-years where the estimated risk is higher, which is the testable prediction of risk-dependent adaptation.

We apply this framework to a panel of 58 cyclone-affected countries from 1980–2019, using the same country set and IBTrACS landfall data as Hsiang and Jina (2014). Our results are as follows. First, the cyclone arrival process is decisively not constant: a pooled specification test rejects the constant Poisson model in favor of a quadratic NHPP with high confidence,

[1] According to the National Oceanic and Atmospheric Administration (NOAA), the United States alone has experienced 391 billion-dollar disasters since 1980, with CPI-adjusted losses totaling $2.75 trillion.

and the quadratic specification dominates the linear trend for every cyclone-affected country in the sample. The intensity path is time-varying and non-monotonic; for many countries, the data are characterized by clusters of arrivals followed by lulls, exactly the kind of pattern a quadratic trend can capture but a linear one cannot. Second, the panel regression delivers a positive, statistically significant interaction coefficient: cyclones do less damage to growth (net of previous year's investment rate and depreciation rate) in country-years when the estimated risk is high. Quantitatively, the baseline cyclone shock lowers per capita growth by roughly one percentage point in a typical country-year, but countries operating at materially elevated risk experience cyclone damages that are smaller by about ten percent of this baseline—a meaningful adaptation offset given that the average cyclone-affected country in our sample is hit roughly once every three to four years.

The bulk of the analysis is devoted to showing that this finding is robust to a series of identification concerns that the recent empirical literature has emphasized. We show that the interaction coefficient is unchanged when we add lagged disaster indicators and country-specific linear trends, ruling out delayed-damage and smooth-trend explanations. We construct placebo distributions by randomly reassigning either the cyclone series or the risk series across countries and show that the actual estimate lies in the upper tail of both.

A central concern in any panel-based identification of adaptation is that the baseline specification is a two-way fixed-effects regression with a continuous moderator. The recent econometrics literature on difference-in-differences with heterogeneous treatment effects (de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021; Borusyak, Jaravel, and Spiess, 2024) has shown that two-way fixed-effects coefficients in such settings can be weighted averages across units with potentially negative or oddly-concentrated weights, so the coefficient need not equal the underlying average treatment effect. The concern is particularly relevant for continuous moderators of the kind we use (Callaway, Goodman-Bacon, and Sant'Anna, 2024).

To address this, we estimate the interaction nonparametrically: sorting country-years into quantile bins of the estimated risk and computing the cyclone damage within each bin yields a smoothly increasing dose-response curve whose slope coincides closely with the panel point estimate. This binned dose-response estimator follows the approach of Callaway, Goodman-Bacon, and Sant'Anna (2024) to continuous-treatment difference-in-differences and avoids the weighting issues at the center of the two-way fixed-effects critique by estimating each bin's treatment effect from a clean within-country comparison. The slope we recover from the nonparametric estimator is in the same direction and similar magnitude as the panel coefficient. Throughout, we use a two-year moving block bootstrap within country to compute standard errors (Politis and Romano, 1994), which we show is the appropriate inference procedure given the short-horizon autocorrelation structure of the residuals; the two-way clustered standard errors used in the prior literature are, by our reckoning, over-conservative by a factor of nearly three in our setting.

We then use the estimated parameters to compute the cumulative welfare value of risk-dependent adaptation over 1980–2019. The thought experiment compares each country's actual 2019 real GDP per capita to the counterfactual path that would have obtained had cyclones inflicted their full unconditional damage in every disaster year, with no risk-dependent offset. Across the cyclone-affected sample, the average country is roughly half a percentage point richer in 2019 because of the risk-dependent adaptation channel, but the cumulative gain reaches two to three percent of real GDP per capita for high-exposure economies whose risk has trended upward over the sample period—such as India, China, Costa Rica, the Dominican Republic, Panama, Thailand, and Bangladesh. These are conservative estimates: the panel coefficient pools across countries and the risk measure captures only deviation from each country's 1980 baseline. The economically more consequential exercise is the forward-looking projection. Combining our estimated interaction coefficient with NHPP-projected risk paths under standard climate-economy growth scenarios shows that the adaptation gain scales sharply with climate risk over time.

Related literature. Our paper is most closely related to the empirical literature on the value of adaptation to climate-related disasters. Hsiang and Jina (2014) and Bakkensen and Mendelsohn (2016) document a robust cross-country pattern: countries with more cumulative cyclone exposure suffer smaller growth losses per cyclone, a result they obtain by augmenting a baseline panel regression with an interaction between the cyclone indicator and a country's historical landfall count n_i. Closely related cross-sectional and temperature-based adaptation findings appear in Gourio and Fries (2020), Carleton, Jina, Delgado, Greenstone, Houser, Hsiang, Hultgren, Kopp, McCusker, Nath, Rising, Rode, Seo, Viaene, Yuan, and Zhang (2022), and Auffhammer (2022); Bilal and Känzig (2024) use global temperature variation. Our framework departs from this literature in two ways. First, our adaptation moderator is the estimated deviation of current risk from the country's baseline, $\tilde{\lambda}_{i,t-1} = \hat{\lambda}_{i,t-1} - \hat{\lambda}_{i,0}$, which varies over time within a country as the underlying disaster process evolves, rather than the cumulative exposure n_i which is fixed by construction. The economic content of our interaction coefficient is therefore "adaptation responds to changes in risk," rather than the static "high-exposure countries adapt more." Second, because $\hat{\lambda}_{it}$ comes from a statistical model of the disaster arrival process, it can be projected forward under alternative climate scenarios—which is what makes the forward-looking adaptation calculation possible. To distinguish these contributions we conduct a direct horse race that includes both the cumulative-exposure interaction $D_{it} \times n_i$ and our time-varying interaction $D_{it} \times \tilde{\lambda}_{i,t-1}$ in the same regression with country and region-by-year fixed effects. Our interaction coefficient survives essentially unchanged at +0.83, while the cumulative-exposure interaction collapses to a statistically insignificant +0.08. The time-varying moderator captures identifying variation that the cumulative-exposure moderator does not.

The paper also speaks to the broader empirical macroeconomics of disasters and weather shocks (Dell, Jones, and Olken, 2012; Deryugina, 2017; Burke, Hsiang, and Miguel, 2015) and to the literature on integrated assessment models of climate change and the social cost of carbon (Nordhaus, 2017; Golosov, Hassler, Krusell, and Tsyvinski, 2014; Barnett, Brock, and Hansen, 2020). For the integrated-assessment literature, our framework provides a forward-projectable estimate of the return to adaptation that can be inserted into climate-economy calibrations alongside projected risk paths under alternative scenarios.

Roadmap. Section 2 describes the data. Section 3 sets up the theoretical motivation and the two-stage estimation strategy. Section 4 reports the first-stage NHPP estimates and the model-selection tests. Section 5 reports the second-stage panel regression and the identification battery (lagged disasters and country trends, placebo tests, horse race against the cumulative-exposure moderator, TWFE-robust binned dose-response). Section 6 computes the in-sample and forward-looking counterfactuals. Section 7 concludes.

2 Data

We combine country-year cyclone landfall data with macroeconomic data for 1980–2019 and use the resulting panel both to estimate country-specific cyclone arrival rates (Section 4) and to identify risk-dependent adaptation in the growth response to landfalls (Section 5).

Tropical cyclones. Cyclone landfall data come from the International Best Track Archive for Climate Stewardship (IBTrACS), the most complete global database of tropical cyclone observations. We define the binary landfall indicator $D_{it} \equiv \text{Landfall}_{i,t}$ as equal to one if and only if country i experienced at least one cyclone landfall of tropical-storm intensity or higher in year t. The country set follows Hsiang and Jina (2014) with two adjustments: we exclude Taiwan, for which World Bank GDP series are unavailable, and for very large countries (the United States, China, Canada) we restrict the landfall definition to the eastern regions that are cyclone-prone, to avoid diluting landfall signal with vast unaffected interiors.[2]

Across the panel, 58 countries experience at least one cyclone landfall in 1980–2019[3], yielding 2,320 country-year observations and 685 landfall events—an unconditional arrival rate of 0.295 per country-year (one landfall roughly every 3.4 years among cyclone-affected countries). The arrival rate varies sharply across the four cyclone basins that we use throughout the paper: East Asia and the Pacific basin (Region 2) sees the highest unconditional intensity at 0.419 per country-year, while the African/South-American basin (Region 4) sees the lowest at 0.056 per country-year. Table 1 reports the regional breakdown, including non-cyclone countries within each basin, so that the disaster-arrival intensity λ in column (3) is interpretable as the basin-wide cyclone exposure rate over the sample.

Table 1: Cyclone landfall frequency by region, 1980–2019

This table reports cyclone landfall counts by region for all countries with a region assignment in the IBTrACS-linked dataset. The four regions correspond to standard cyclone basins: Region 1 = North Atlantic + Western Europe (including North America and the Caribbean); Region 2 = East Asia + Pacific (including Oceania); Region 3 = Indian Ocean + Middle East/North Africa + South Asia; Region 4 = Africa + South America. Column (3) is the unconditional disaster arrival intensity $\hat{\lambda} = (2)/(1)$, the baseline cyclone landfall rate per country-year used as the starting point for the time-varying NHPP estimates in Section 4. The Total row covers the 108 countries with a region assignment; an additional 10 cyclone-affected countries lack region assignments and enter the regression sample only when region is not required.

Region (cyclone basin)	Country-year obs.	Landfalls	$\hat{\lambda}$
R1: N. Atlantic + W. Europe	1,240	177	0.143
R2: E. Asia + Pacific	520	218	0.419
R3: Indian Ocean + MENA + S. Asia	600	47	0.078
R4: Africa + S. America	1,960	110	0.056
Total (108 countries)	4,320	552	0.128

Economic data. The key macroeconomic variable is real GDP per capita, Y_{it}, taken from the World Bank World Development Indicators in constant 2015 US dollars. We work throughout with the log growth rate $g_{it} \equiv \log(Y_{it}/Y_{i,t-1})$, expressed in percentage points. We also use country-level investment and depreciation flows from Penn World Table 10.0. Investment is current-period gross fixed capital formation, I_{it}, scaled by lagged output to give an investment-to-output ratio $i_{it} \equiv I_{it}/Y_{i,t-1}$; the depreciation variable is the current-period depreciation flow

[2] We use pre-1980 cyclone observations, 1960–1979, to inform initial-condition estimates for the country-specific arrival-rate models in Section 4.

[3] Cross-country data on cyclones before 1980 are sparse. We end the series in 2019 before Covid-19.

D^{δ}_{it} scaled by lagged output, $\delta_{it} \equiv D^{\delta}_{it}/Y_{i,t-1}$.[4] Both i_{it} and δ_{it} are therefore on the same dimensional scale as the log growth rate: each is expressed as a fraction of lagged output.

Motivated by Equation (3), our second-stage regressions use growth net of the investment-minus-depreciation contribution:

$$g^n_{it} \;\equiv\; g_{it} - (i_{it} - \delta_{it}) \;=\; \log\left(\frac{Y_{it}}{Y_{i,t-1}}\right) \;-\; \frac{I_{it} - D^{\delta}_{it}}{Y_{i,t-1}}. \tag{1}$$

The decomposition is dimensionally consistent: for small changes, $\log(Y_{it}/Y_{i,t-1}) \approx (Y_{it} - Y_{i,t-1})/Y_{i,t-1}$ is itself a fraction of lagged output, so subtracting $(I_{it} - D^{\delta}_{it})/Y_{i,t-1}$—the net change in capital as a fraction of lagged output—leaves a residual on the same scale. This is the empirical analog of the disaster-and-noise residual $F(\lambda)dJ + \sigma dW$ in the continuous-time decomposition $g_t = (i(\lambda) - \delta) + F(\lambda)dJ_t + \sigma dW_t$ of Hong, Wang, and Yang (2023) (Equation 3): netting out the smooth investment-minus-depreciation term leaves the part of growth that responds to disasters and to idiosyncratic shocks.

This construction has two practical advantages over working with raw growth directly. First, the smooth $(i_{it} - \delta_{it})$ component absorbs the medium-frequency variation in growth that is driven by capital accumulation, leaving a residual g^n_{it} with standard deviation roughly one tenth that of g_{it} (0.46 vs 4.59 percentage points in our cyclone sample). The disaster signal is detectable in g^n_{it} but is buried in the much noisier raw growth series. Second, the unit-coefficient restriction on $(i_{it} - \delta_{it})$ is the empirical version of the model restriction: the deterministic component of output growth equals the deterministic component of capital growth (net of depreciation), which is exactly Equation (3) written without the disaster term.[5]

Table 2 reports summary statistics across the 58-country cyclone sample. The mean log growth rate of real GDP per capita is 3.4 percentage points with a standard deviation of 4.6 percentage points. The mean investment ratio i_{it} is 0.23 with a standard deviation of 0.09; the mean depreciation ratio δ_{it} is 0.21 with a standard deviation of 0.09. The net-investment margin $i_{it} - \delta_{it}$ thus averages about 1–2 percentage points of lagged output in the cyclone sample. The mean growth-net residual g^n_{it} is -0.21 percentage points with a standard deviation of

[4] Note that the depreciation ratio δ_{it} defined this way is the depreciation flow as a fraction of lagged output, not the depreciation rate of capital. With a capital-output ratio of $K/Y \approx 3$–4 and a depreciation rate of capital around 5–10%, δ_{it} typically takes values in the 0.15–0.25 range in our panel.

[5] An alternative that has been suggested is to use raw g_{it} on the left-hand side and to include i_{it} and δ_{it} as separate right-hand-side controls with freely-estimated coefficients. Under the HWY decomposition this is equivalent to our specification only if those coefficients equal -1 and $+1$, respectively. With freely-estimated coefficients, the unit-coefficient restriction is not imposed, the cyclone-related variation is partly absorbed into the (i, δ) controls, and the disaster effect on growth is no longer cleanly identified. We discuss this further when introducing the baseline regression in Section 5.

0.46 percentage points, consistent with g^n absorbing the average investment-minus-depreciation contribution and leaving a near-zero residual to be explained by disaster realizations.

Table 2: Summary statistics of economic variables, cyclone sample, 1980–2019

This table reports summary statistics of the key macroeconomic variables across the 58-country cyclone-affected sample. $g_{it} = 100 \cdot \log(Y_{it}/Y_{i,t-1})$ is the log growth rate of real GDP per capita in percentage points (pp; World Bank WDI, constant 2015 US$). The investment ratio $i_{it} = I_{it}/Y_{i,t-1}$ is current-period gross fixed capital formation scaled by lagged output (PWT 10.0); the depreciation ratio $\delta_{it} = D^{\delta}_{it}/Y_{i,t-1}$ is the current-period depreciation flow scaled by lagged output (PWT 10.0). Both i_{it} and δ_{it} are decimal fractions of lagged output. $g^n_{it} = g_{it} - 100 \cdot (i_{it} - \delta_{it})$ is growth net of the change in capital stock per lagged output, in pp (see Equation (1); all three terms on the same pp scale). N is the number of non-missing country-year observations for each variable.

	N	Mean	S.D.	Median	P10	P90
g_{it} (pp)	2,279	3.37	4.59	3.61	−0.99	7.99
i_{it} (fraction of $Y_{i,t-1}$)	2,262	0.23	0.09	0.22	0.13	0.34
δ_{it} (fraction of $Y_{i,t-1}$)	2,262	0.21	0.09	0.20	0.11	0.31
g^n_{it} (pp)	2,262	−0.21	0.46	−0.02	−1.14	0.06

3 Extreme-Weather Damages: Theory and Estimation

A widely-used approach to estimating the value of adaptation to disasters such as tropical cyclones (Bakkensen and Mendelsohn (2016), Hsiang and Jina (2014)) is to measure how economic damage conditional on a disaster arrival significantly declines cross-sectionally with the historical experience of a locale with disasters (see, e.g., Dell, Jones, and Olken (2014) for a review).[6] They augment a baseline panel regression of GDP growth on the disaster indicator D_{it} with country and year fixed effects by estimating the following specification:

$$g_{it} = \phi D_{it} + (\psi \cdot n_i) \times D_{it} + u_i + v_{jt} + \varepsilon_{it}, \tag{2}$$

where n_i is the variable measuring the number of extreme weather events a country i has experienced historically, u_i is a country fixed effect, and v_{jt} is a region-by-time fixed effect. Adaptation studies typically find that ψ is positive, so that countries with more extreme-weather arrivals historically (measured by high n_i values) experience smaller damage for a given event D_{it} ($\phi + \psi n_i$ is less negative than ϕ alone). Thus in this reduced-form model (2), the coefficient ϕ captures the conditional damage without adaptation and the coefficient ψ captures the adaptation effect.

[6] This time-invariant approach, which goes back to earlier work on identifying the impact of temperature on agricultural yields (Deschênes and Greenstone (2007) and Schlenker and Roberts (2009)), relies on location and time fixed effects to address unobserved heterogeneities in the panel data.

This time-invariant approach is a special case of our setting in which the level of adaptation is fixed at the country's prior risk level. Closely related cross-sectional adaptation findings appear in Auffhammer (2022), Gourio and Fries (2020), and Carleton, Jina, Delgado, Greenstone, Houser, Hsiang, Hultgren, Kopp, McCusker, Nath, Rising, Rode, Seo, Viaene, Yuan, and Zhang (2022).[7]

Though simple and intuitive, this regression has a substantive limitation: the moderator n_i is fixed by construction, so the specification is silent on whether and how a country adapts as its underlying disaster risk changes over time. This omission is increasingly hard to defend in a setting where climate change is projected to push disaster intensities outside their historical ranges, and where the policy-relevant question is precisely how the marginal damage of a cyclone responds to a shifting risk environment. The recent theoretical literature on optimal adaptation, building on Hong, Wang, and Yang (2023), addresses exactly this question by allowing the per-disaster damage to depend on the contemporaneous arrival rate.[8]

We suppress the country subscript i in describing the theoretical model. In the continuous-time framework of Hong, Wang, and Yang (2023), disasters arrive according to a Poisson process with intensity λ_t. The model predicts that agents adapt to disasters, in the sense that economic outcomes (such as economic growth) depend on the state variable λ_t. For example, output growth g_t (adjusted for investment and depreciation) takes the form

$$g_t - (i(\lambda_{t-}) - \delta)dt = F(\lambda_{t-})dJ_t + \sigma dW_t. \tag{3}$$

where i is the investment-to-capital ratio, δ is the depreciation rate, $F(\lambda_{t-})$ is an adaptation function that mediates damage to growth caused by the jump dJ_t and depends on λ_t, and dW_t is an idiosyncratic Brownian shock. The output response to disasters ($dJ_t = 1$) is state-dependent with two features: it is time varying (to the extent that λ_{t-} varies with t) and possibly non-linear (to the extent that $F(\cdot)$ is non-linear). As λ_{t-} rises, there is more adaptation spending which lowers economic damage.

The empirical approach in the rest of the paper is motivated directly by this theoretical structure. We estimate the model in two stages. First, we estimate the country-specific arrival rate λ_{it} as a nonhomogeneous Poisson process whose intensity follows a deterministic time trend; this is the first stage and the object of Section 3.1 below. Second, we relate observed growth to disaster realizations and the estimated arrival-rate path $\hat{\lambda}_{it}$ in a panel regression that delivers

[7] An exception is Bilal and Känzig (2024), who use aggregate (global) time-series variation to measure economic damages from changing average temperatures.

[8] Empirically, Barreca, Clay, Deschenes, Greenstone, and Shapiro (2016) document a declining relationship over time in the temperature-mortality relationship due to increasing adaptation in the form of air conditioning.

the slope of the adaptation function $F(\cdot)$; this is the second stage and the object of Section 3.2 below.

3.1 Estimating Time-Varying Extreme-Weather Risk

We model each country's cyclone arrival process as a nonhomogeneous Poisson process (NHPP). If $\{N_t : t \geq 0\}$ is the counting process that records the number of events up to time t, the deterministic time-trend NHPP is defined through an intensity function λ_t such that the expected increment in N_t over an infinitesimal interval $(t, t+dt]$ is $\lambda_t \, dt$ and increments are independent.

The intensity function is typically parameterized to ensure positivity through a log-linear model:

$$\lambda_t = \exp(x_t' \theta), \tag{4}$$

where x_t is a vector of covariates (possibly including deterministic functions of time) and θ is the parameter vector to be estimated. We consider two models for λ_t, both of which are deterministic functions of time:

$$\lambda_t = \exp(\theta_0 + \theta_1 t) \quad \text{the trend model (linear time trend)}, \tag{5}$$

$$\lambda_t = \exp(\theta_0 + \theta_1 t + \theta_2 t^2) \quad \text{the quadratic } t \text{ model (quadratic time trend)}. \tag{6}$$

For our observations of 0/1 arrival indicators $D_t; t = 1, \ldots, T$, the likelihood function of the time-trend model is (see Daley and Vere-Jones (2003) and Kutoyants (1998) for a reference):

$$L(\theta) = \exp\left(- \sum_{t=1}^{T} \lambda_t \right) \prod_{t=1}^{T} \lambda_t^{D_t}, \tag{7}$$

where λ_t is either the trend model specified in (5) or the quadratic model specified in (6). The log-likelihood function is therefore:[9]

$$\mathcal{L}(\theta) = - \sum_{t=1}^{T} \lambda_t + \sum_{t=1}^{T} D_t \log\big(\lambda_t\big). \tag{8}$$

The parameter vector θ of our time-trend models can then be directly estimated by maximum likelihood using the log-likelihood function (8).

[9] See Equation (7.1.2), Chapter 7.1 of Daley and Vere-Jones (2003) for the same log-likelihood function.

3.2 Estimating Risk-Dependent Damages

Equation (3) suggests a flexible alternative to the linear model

$$g_{it} - (i_{it} - \delta_{it}) = \mu_i + F(\lambda_{it-1})D_{it} + \varepsilon_{it}, \tag{9}$$

where $F(\lambda_{it-1})$ is the adaptation-induced damage function. Rather than modeling the growth rate g_{it}, we work with GDP growth net of the investment-minus-depreciation contribution, $g^n_{it} = g_{it} - (i_{it} - \delta_{it})$, where (as defined in Section 2) $g_{it} = \log(Y_{it}/Y_{i,t-1})$ is the log growth of per-capita output and $i_{it} = I_{it}/Y_{i,t-1}$, $\delta_{it} = D^{\delta}_{it}/Y_{i,t-1}$ are current-period investment and depreciation flows scaled by lagged output. All three terms are on the same fraction-of-lagged-output scale, so the netting is dimensionally consistent.[10] Note that all variables in the regression are country-specific.

The standard constant-coefficient linear regression in the literature obtains as a special case when $F(\lambda_{it-1})$ is a constant $\zeta_{1i} = F(\lambda_{i0})$, i.e. when there is no time trend in arrivals. Then

$$g^n_{it} = \mu_i + \zeta_{1i} D_{it} + \varepsilon_{it}. \tag{10}$$

The growth of an economy depends on investment (based on λ_{it-1}). We expect the constant term μ_i to be zero (see also the discussion in Section 2).

A non-parametric specification with $\hat{\lambda}_{it}$ and D_{it} as inputs would estimate the conditional damage function flexibly but would not isolate the structural object of interest: the slope of $F(\cdot)$ at the country's baseline risk, which is what governs the adaptation response to a marginal change in risk. We therefore impose a first-order Taylor expansion of $F(\lambda_{it-1})$ around λ_{i0}, which delivers a linear-in-parameters specification while preserving the structural interpretation of the slope coefficient:

$$\begin{aligned} g^n_{it} &= \mu_i + \underbrace{F(\lambda_{i0})}_{\substack{\text{damage at} \\ \text{prior adaptation}}} D_{it} + \underbrace{F'(\lambda_{i0})(\lambda_{it-1} - \lambda_{i0})}_{\substack{\text{return to risk-dependent} \\ \text{adaptation}}} D_{it} + \varepsilon_{it} \\ &= \mu_i + \beta_{1i} D_{it} + \beta_{2i} \tilde{\lambda}_{it-1} D_{it} + \varepsilon_{it} \\ &= \mu_i + \chi_{it} + \varepsilon_{it} \end{aligned} \tag{11}$$

[10] An alternative would be to use raw growth g_{it} on the left-hand side and include i_{it} and δ_{it} as separate right-hand-side controls. Under the model decomposition $g_t = (i(\lambda) - \delta) + F(\lambda)D + \sigma\, dW$, that specification is equivalent to ours only if the coefficient on $(i_{it} - \delta_{it})$ is restricted to one. With the coefficient estimated freely, the cyclone signal is buried in the much larger growth-rate noise: in our cyclone-affected sample, g_{it} has a standard deviation roughly 10 times that of g^n_{it} (4.59 vs. 0.46 percentage points), and the model-implied unit-coefficient restriction is what makes the disaster effect identifiable. We therefore work with g^n_{it} throughout.

where $\tilde{\lambda}_{it-1} = (\lambda_{it-1} - \lambda_{i0})$ and $\chi_{it} = \beta_{1i} D_{it} + \beta_{2i} \tilde{\lambda}_{it-1} D_{it}$ is the total damage in the presence of risk-dependent adaptation. The $F(\lambda_{i0})$ term is economic damage with adaptation fixed at the country's baseline risk λ_{i0}. The $F'(\lambda_{i0})\tilde{\lambda}_{it-1}$ term is the return to risk-dependent adaptation induced by changes in the time-varying arrival rate, net of the cost of adaptation to society.

Note that $F(\lambda_{i0})$ and $F'(\lambda_{i0})$ suffice for identifying the first-order effects of D_t, being

$$E[g^n_{it}|D_{it} = 1] = \underbrace{F(\lambda_{i0})}_{\beta_{1i}<0} + \underbrace{F'(\lambda_{i0})}_{\beta_{2i}>0} \tilde{\lambda}_{it-1} \tag{12}$$

Equation (11) nests the linear model (10) in which F is constant (no risk-dependent adaptation) and $(\lambda_{it-1} - \lambda_{i0}) = 0$ (no time trend). Equation (11) also nests (2) in which D_{it} is not interacted with λ_{it-1}. Omitting dependence on λ_{i0} and λ_{it-1} could bias the estimates of economic damages.

4 Time-Varying Disaster Risks

We estimate the country-specific cyclone arrival rate by maximum likelihood under three nested NHPP specifications: a Constant rate, a Linear time-trend, and a Quadratic time-trend,

$$\hat{\lambda}^{\text{const}}_{it} = \exp(\hat{\theta}_{0i}), \qquad \hat{\lambda}^{\text{lin}}_{it} = \exp(\hat{\theta}_{0i} + \hat{\theta}_{1i} t), \qquad \hat{\lambda}^{\text{quad}}_{it} = \exp(\hat{\theta}_{0i} + \hat{\theta}_{1i} t + \hat{\theta}_{2i} t^2),$$

each fit to country i's cyclone landfall sequence by maximizing the NHPP log-likelihood in Equation (8). Two questions guide model selection. First, are cyclone arrival rates time-varying at all—can the constant Poisson model be rejected? Second, conditional on time variation, is the linear trend sufficient or do we need the quadratic flexibility?

Pooled Vuong testing. With only 5–15 cyclone arrivals per country over a 40-year sample, country-by-country likelihood-ratio tests have very low power and are individually uninformative for most countries. We therefore lean on a pooled Vuong (1989) test that combines per-period log-likelihood differences across countries. Define $\hat{ll}^M_{it} = -\hat{\lambda}^M_{it} + D_{it} \log \hat{\lambda}^M_{it}$ as the per-period log-likelihood under model $M \in \{\text{const,lin,quad}\}$, and let $v_{it} = \hat{ll}^A_{it} - \hat{ll}^B_{it}$ be the per-period log-likelihood difference between two candidate models A and B. Stacking v_{it} across all N countries and T_i years of data yields $T = \sum_i T_i$ country-year observations. Under the null of equal expected fit, the standardized statistic

$$V = \sqrt{T}\, \bar{v} \,/\, \hat{\sigma}_v \tag{13}$$

is asymptotically standard normal, where $\bar{v}$ and $\hat{\sigma}_v^2$ are the pooled mean and variance of v_{it}. Positive (negative) values of V favor model A (B). We pair this with a country-level sign test: under the null of equal fit, the share of countries for which model A attains a higher log-likelihood than model B should equal 50%, with deviations summarized by a one-sided binomial p-value.

Table 3 reports both diagnostics for all three pairwise comparisons. The constant Poisson model is decisively rejected in favor of the quadratic NHPP: $V = +7.30$ (two-sided $p < 0.0001$), and for 58 out of 58 cyclone-affected countries the quadratic specification attains a strictly higher log-likelihood than the constant. The linear model also dominates the constant ($V = +4.47$, $p < 0.0001$), but the quadratic specification adds meaningful flexibility on top of the linear trend: $V = +5.79$ ($p < 0.0001$), again with a clean sign sweep of 58/58 countries. The takeaway is uniform across the panel: cyclone arrival rates are not constant, and a linear time trend is not flexible enough to capture the time variation that the data reveal.

Table 3: NHPP Model Selection Tests for the Cyclone Arrival Rate

This table reports model-selection tests across the 58 cyclone-affected countries (countries with at least one landfall in 1980–2019). For each country we estimate by maximum likelihood three nested NHPP specifications of the cyclone arrival rate λ_{it}: a Constant model $\lambda_{it} = \exp(\theta_{0i})$, a Linear time-trend model $\lambda_{it} = \exp(\theta_{0i} + \theta_{1i}t)$, and a Quadratic time-trend model $\lambda_{it} = \exp(\theta_{0i} + \theta_{1i}t + \theta_{2i}t^2)$. Per-country likelihood-ratio tests have low power given the small number of landfalls per country (typically 5–15), so we report two pooled diagnostics. The pooled Vuong statistic stacks the per-period log-likelihood differences $v_{it} = \hat{ll}_{it}^A - \hat{ll}_{it}^B$ across all country-years and tests H_0: $E[v_{it}] = 0$; under the null $V = \sqrt{T}\,\bar{v}/\hat{\sigma}_v$ is asymptotically standard normal, with T the total number of country-years. The sign test reports the share of countries for which the more flexible model attains a strictly higher (or equal) log-likelihood; under the null of equal fit each country has a 50% chance of favoring either model, and the one-sided binomial p-value is reported. ***, **, * denote significance at the 0.1%, 1%, and 5% levels.

Comparison (A vs B)	Pooled Vuong V	Two-sided p	Sign: $\hat{ll}^A \geq \hat{ll}^B$
Quadratic vs Constant	$+7.30^{***}$	< 0.0001	58/58 (100%)***
Linear vs Constant	$+4.47^{***}$	< 0.0001	58/58 (100%)***
Quadratic vs Linear	$+5.79^{***}$	< 0.0001	58/58 (100%)***
Countries		58	
Country-years (pooled)		2,320	

Fiji as illustration. Figure 1 illustrates the model-selection result for Fiji's cyclone series. Fiji experienced ten landfall years in 1980–2019, with a cluster in the early 1980s, a second cluster in the 1990s, a long lull from 1996 to 2012, and then a renewed cluster in 2013–2019. The constant model ($\hat{\lambda} = 0.25$) and the linear trend (which is nearly flat for Fiji) both miss the U-shape entirely. The quadratic NHPP traces a declining-then-rising path: it correctly puts higher arrival probability in the early 1980s, drops to its trough during the 1996–2012 lull, and rises sharply again during the late-sample cluster. The per-period log-likelihood gain from quadratic over linear (+1.17 for Fiji alone) is on the order of magnitude of an additional parameter's contribution to fit, and similar U-shaped (or inverted-U) patterns drive the pooled rejection across the panel.

Figure 1: Fitted NHPP arrival rate and actual cyclone landfalls: Fiji

This figure illustrates the three NHPP arrival rate specifications fit to Fiji's cyclone landfall sequence, 1980–2019. Gray bars mark years with a cyclone landfall ($D_t = 1$). The black dotted line is the constant model ($\hat{\lambda} = 0.25$); the orange line is the linear time-trend model; the blue line is the quadratic time-trend model. The quadratic fit captures the U-shaped pattern (high arrival probability in the early 1980s and again post-2013, with a deep lull in the late-1990s and 2000s) that the linear and constant models cannot accommodate.

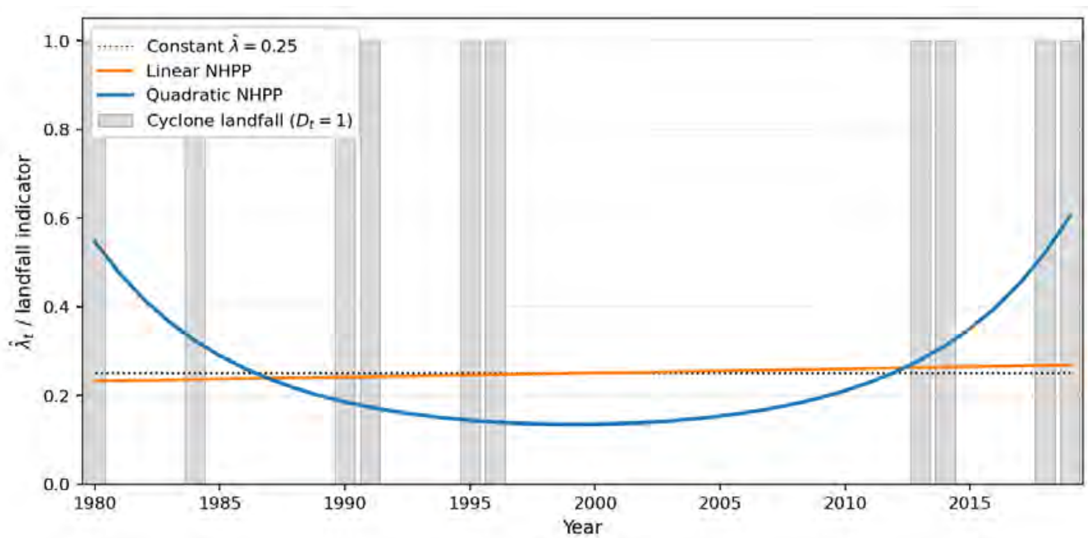

Distribution of fitted arrival rates and deviations. Before turning to the second-stage regression, it is useful to summarize the distribution of the fitted NHPP arrival rates across the panel, since this distribution determines the empirical scale on which the adaptation-channel coefficient β_2 in Section 5 will be interpreted. Table 4 reports the cross-sectional distribution of three quantities: the country-specific 1980 baseline $\hat{\lambda}_{i,0}$ (one value per country), the year-t arrival rate $\hat{\lambda}_{it}$ (one value per country-year), and the deviation from baseline $\tilde{\lambda}_{it} \equiv \hat{\lambda}_{it} - \hat{\lambda}_{i,0}$ (one value per country-year). All three quantities are unitless probabilities or rates in $[0, 1]$, and we report them as decimal fractions throughout. The 1980 baseline $\hat{\lambda}_{i,0}$ has substantial cross-country dispersion—mean 0.31, median 0.18, range $[0.00, 1.00]$—reflecting the differences in cyclone exposure across the cyclone basins documented in Table 1. The pooled arrival rate $\hat{\lambda}_{it}$ has a slightly higher mean (0.36) and similar dispersion, indicating that on average the panel has experienced some upward drift in arrival rates relative to its 1980 baseline. The key object for our analysis is the deviation $\tilde{\lambda}_{it}$, which has mean +0.043, standard deviation 0.10, and ranges from -0.15 (countries whose current risk has fallen below the 1980 baseline) to +0.37 (countries whose current risk has risen sharply above the 1980 baseline). This last quantity sets the empirical scale for the adaptation coefficient β_2 in the second stage: a coefficient of +0.86 in percentage points per unit of $\tilde{\lambda}$ translates to a +0.086 pp adaptation offset at a one-standard-deviation increase in $\tilde{\lambda}$, and a +0.32 pp offset at the upper-tail country-year value of $\tilde{\lambda} \approx 0.37$.

Table 4: Distribution of estimated NHPP arrival rates across cyclone-affected countries, 1980–2019

This table reports the cross-sectional distribution of the quadratic NHPP arrival-rate estimates used in the second-stage regression. $\hat{\lambda}_{i,0}$ is country i's 1980 baseline arrival rate (one value per country, $N = 58$); $\hat{\lambda}_{it}$ is the year-t arrival rate from the country-specific NHPP fit (one value per country-year); $\tilde{\lambda}_{it} = \hat{\lambda}_{it} - \hat{\lambda}_{i,0}$ is the deviation of current risk from the country's 1980 baseline (one value per country-year). All quantities are unitless probabilities or rates in $[0, 1]$. The deviation $\tilde{\lambda}_{it}$ has SD ≈ 0.10 across the panel and ranges from -0.15 (countries whose risk has fallen below the 1980 baseline) to +0.37 (countries whose risk has risen sharply above the 1980 baseline).

	N	Mean	S.D.	Median	P10	P25	P75	P90
$\hat{\lambda}_{i,0}$ (1980 baseline)	58	+0.310	0.345	+0.184	+0.003	+0.049	+0.424	+0.976
$\hat{\lambda}_{it}$ (year-t intensity)	2,320	+0.355	0.370	+0.186	+0.000	+0.024	+0.634	+0.981
$\tilde{\lambda}_{it}$ (deviation from baseline)	2,320	+0.043	0.102	+0.004	-0.051	-0.026	+0.085	+0.232

5 Return to Risk-Dependent Adaptation

Having concluded our discussion of the first-stage country risk dynamics, we turn to our second-stage estimates. The standard approach in the literature is to estimate average disaster damages by a pooled panel regression with country and year fixed effects. We differ by allowing the damage function to vary with $\tilde{\lambda}_{i,t-1} = \hat{\lambda}_{i,t-1} - \hat{\lambda}_{i,0}$, the lagged deviation of the time-varying cyclone arrival rate from its country baseline.

5.1 Baseline Panel Regression

We estimate the following panel regression of real GDP per capita growth (net of investment less depreciation) on all countries with at least one cyclone landfall in 1980–2019:

$$g^n_{it} = \mu_i + \nu_{r(i),t} + \beta_1 D_{it} + \beta_2 \, \tilde{\lambda}_{i,t-1} \, D_{it} + \varepsilon_{it}, \tag{14}$$

where μ_i is a country fixed effect and $\nu_{r(i),t}$ is a region-by-year fixed effect (with $r(i)$ denoting the region to which country i belongs). We use a four-region classification along cyclone-basin lines: (1) North Atlantic and Europe, (2) East Asia and Pacific, (3) Indian Ocean and Middle East, and (4) Africa and South America. Including region $\times$ year fixed effects (as opposed to year fixed effects alone) absorbs any region-specific year shocks—most importantly, the shared component of climate-change-driven trends in cyclone risk within a region. As a consequence, β_2 is identified entirely by within-country, within-region-year variation: it captures whether country i's GDP growth in a cyclone year deviates more or less from its region's growth that year, depending on country i's position relative to its region's risk trajectory. The dependent variable g^n_{it} is in percentage points (Section 2 and Equation (1)); D_{it} is binary; and $\tilde{\lambda}_{i,t-1}$ is a unitless arrival-rate deviation in $[-0.15, +0.37]$ (Table 4). Both coefficients $\hat{\beta}_1$ and $\hat{\beta}_2$ are therefore in percentage points: $\hat{\beta}_1$ is the change in g^n in pp from a $D = 0 \to D = 1$ shift, and $\hat{\beta}_2$ is the change in g^n in pp per unit of $\tilde{\lambda}$ within a disaster year. β_1 captures the average direct damage from a cyclone; β_2 captures the extent to which damage is attenuated when prior risk is high—our measure of risk-dependent adaptation. Standard errors are computed via a two-year moving block bootstrap within country, which preserves short-run residual autocorrelation while not over-conservatively treating all within-country observations as a single cluster.

Table 5 reports the baseline estimates on the sample of 48 countries with at least one cyclone landfall and non-missing region assignment. Column (1) presents the average direct effect of a cyclone landfall: a cyclone reduces real GDP growth by approximately one percentage point ($\hat{\beta}_1 = -0.93$, $t_{\text{MBB}} = -41.1$). Column (2) adds the interaction with prior risk $\tilde{\lambda}_{i,t-1}$. The direct effect tightens slightly ($\hat{\beta}_1 = -0.98$), and the interaction coefficient is positive and highly significant ($\hat{\beta}_2 = +0.82$, $t_{\text{MBB}} = +4.48$, $p < 0.001$). The positive sign on β_2 is consistent with risk-dependent adaptation: countries facing periods of elevated cyclone risk—above and beyond their region's contemporaneous trend—exhibit attenuated GDP responses when cyclones strike. Column (3) adds the standalone lagged arrival rate $\tilde{\lambda}_{i,t-1}$ as an additional control, γ. This regressor would capture any direct effect of climate risk on GDP growth outside of disaster years that is not absorbed by the region $\times$ year fixed effects. The estimated $\hat{\gamma} = -0.10$ is small in magnitude and statistically insignificant ($t_{\text{MBB}} = -1.53$, $p = 0.13$), while $\hat{\beta}_2$ is essentially unchanged at $+0.86$ and remains highly significant ($t_{\text{MBB}} = +4.67$, $p < 0.001$). The risk-dependent attenuation is therefore not driven by a confounding direct effect of climate risk on growth; it operates specifically through the disaster-contingent interaction.

The economic magnitude of the adaptation interaction is substantial. Within the regression sample, the lagged risk variable $\tilde{\lambda}_{i,t-1}$ has a standard deviation of 0.10. A one-standard-deviation increase in prior risk therefore offsets the direct cyclone damage by $\hat{\beta}_2 \times 0.10 \approx 0.085$ percentage

Table 5: Baseline Panel Regressions: Risk-Dependent Cyclone Damages

This table reports panel regressions of real GDP per capita growth (net of investment less depreciation) on the cyclone indicator D_{it}, its interaction with the lagged time-varying arrival rate $\tilde{\lambda}_{i,t-1} = \hat{\lambda}_{i,t-1} - \hat{\lambda}_{i,0}$, and in Column (3) the lagged arrival rate $\tilde{\lambda}_{i,t-1}$ on its own. $\hat{\lambda}_{it}$ is estimated from a country-specific quadratic time trend (NHPP). The sample includes all countries with at least one cyclone landfall in 1980-2019. Country fixed effects and region $\times$ year fixed effects are included throughout (the region variable groups countries into 4 cyclone-basin regions, so region $\times$ year FE absorb any shared regional year shocks). Standard errors (in parentheses) are computed via a 2-year moving block bootstrap within country (B=1000 replications). ***, **, * denote significance at the 1%, 5%, and 10% levels.

	(1)	(2)	(3)
β_1 (cyclone D_{it})	-0.9328***	-0.9798***	-0.9839***
	(0.0227)	(0.0278)	(0.0279)
γ ($\tilde{\lambda}_{i,t-1}$)			-0.0984
			(0.0643)
β_2 ($D_{it} \times \tilde{\lambda}_{i,t-1}$)		+0.8222***	+0.8590***
		(0.1836)	(0.1839)
Country FE	Yes	Yes	Yes
Region $\times$ Year FE	Yes	Yes	Yes
Observations	1,872	1,872	1,872
Countries	48	48	48

points, or about 8.7% of $|\hat{\beta}_1|$. A two-standard-deviation shift offsets roughly 17% of the baseline damage. At the upper end of the sample's risk distribution ($\tilde{\lambda} \approx +0.37$, observed for cyclone-active country-years such as Pakistan in the late 1980s and Senegal in the early 1980s), the net cyclone effect on GDP growth shrinks to approximately -0.66 percentage points—a 33% reduction in damage relative to the average.

Figure 2 visualizes the partial-correlation interpretation of $\hat{\beta}_2$. By the Frisch–Waugh–Lovell theorem, the OLS slope of the residual of g^n_{it} on the residual of $D_{it} \times \tilde{\lambda}_{i,t-1}$—both residualized on country fixed effects, region $\times$ year fixed effects, D_{it}, and $\tilde{\lambda}_{i,t-1}$—equals $\hat{\beta}_2$ from the full panel regression in Column (3) of Table 5. The figure shows the residual cloud (grey), a 20-bin equal-frequency binscatter of the residuals (blue, with 95% confidence intervals), and the fitted line. The binscatter markers track the OLS fit closely and the slope coincides exactly with the panel estimate $\hat{\beta}_2 = +0.859$: the upward partial-correlation pattern is visible in the data conditional on the controls, and is not the artifact of a single influential observation or a non-linear functional form.

Standard errors via a 2-year moving block bootstrap. The standard errors reported in parentheses below the coefficients in Table 5 are obtained from a moving block bootstrap (MBB) within country (Politis and Romano (1994)). For each bootstrap replication, we resample, within each country, blocks of two consecutive years of observations with replacement, concatenate the blocks into a new country time series of the same length as the original, and re-estimate the panel regression on the resampled data. The bootstrap standard error of a coefficient is the standard deviation of its estimates across $B = 1{,}000$ such replications. The block size of two years is conservative relative to the empirical autocorrelation structure of the regression residuals, which is statistically negligible at all lags: across the 48 countries in the regression sample, the pooled within-country residual autocorrelation coefficient is -0.01 at lag 1 ($p = 0.56$), -0.05 at lag 2 ($p = 0.12$), and -0.03 at lag 3 ($p = 0.38$), and country-by-country tests reject the null of no first-order autocorrelation for only 8% of countries—i.e., below the nominal Type I error rate.

The MBB is preferred over conventional two-way cluster-robust standard errors (clustered by country and year) for two reasons. First, two-way clustering implicitly allows residuals within a country to be correlated across arbitrarily long time intervals—treating an observation in 1980 as potentially correlated with one in 2019. The serial correlation tests above show this is empirically not the case: residual dependence within country is short-lived, if present at all. Imposing potentially long-range within-country dependence therefore produces standard errors that are far larger than the data warrant. Second, with only 40 years in the time-cluster dimension, the asymptotic theory underlying conventional cluster-robust inference is known to

Figure 2: Partial-correlation view of $\hat{\beta}_2$: residualized scatter and binscatter

This figure plots the Frisch–Waugh–Lovell residualization of the baseline panel regression (Table 5, Column (3)). The horizontal axis is the residual of $D_{it} \times \tilde{\lambda}_{i,t-1}$ from a regression on country fixed effects, region $\times$ year fixed effects, D_{it}, and $\tilde{\lambda}_{i,t-1}$. The vertical axis is the residual of g^n_{it} from the same controls. Grey points are the underlying residual cloud (a random sample for legibility). Blue markers are the binscatter at 20 equal-frequency bins of the horizontal-axis residual, with 95% confidence intervals based on within-bin standard deviations. The red line is the OLS fit; its slope of +0.859 equals the panel coefficient $\hat{\beta}_2$ exactly by the Frisch–Waugh–Lovell theorem.

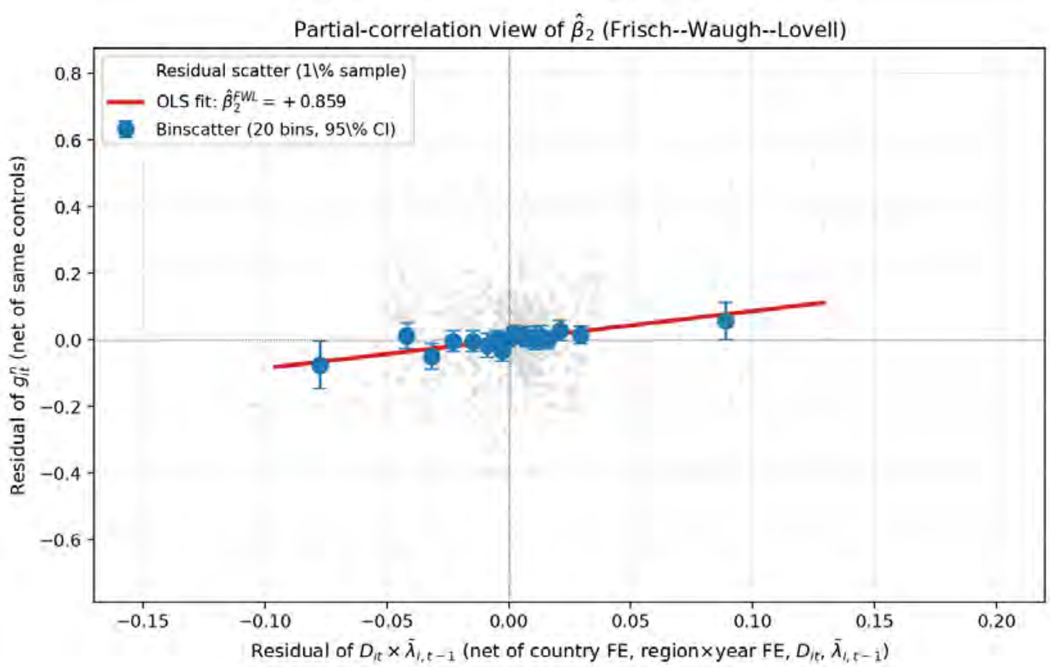

perform poorly: cluster-robust standard errors can be biased downward when the number of clusters in a dimension is small, leading to over-rejection of the null (MacKinnon and Webb (2018)). Bootstrap-based alternatives such as the MBB are recommended in such settings because they avoid the small-cluster asymptotic approximation entirely.

Table 6 compares the MBB standard errors for $\hat{\beta}_1$ and $\hat{\beta}_2$ in our baseline specification with three conventional alternatives: heteroskedasticity-robust, country-clustered, and two-way (country and year) clustered. The coefficient estimates are identical across methods; only the standard errors differ. Two patterns are notable. First, the MBB and the heteroskedasticity-robust standard errors are nearly identical for $\hat{\beta}_2$ (0.184 vs. 0.209). The two methods coincide whenever residuals carry no within-country serial correlation, which the residual autocorrelation diagnostics above confirm to be the case here. Second, the country-clustered and two-way-clustered standard errors are roughly 2.7 times larger than the MBB—and substantially overstate the uncertainty in $\hat{\beta}_2$. Cluster-robust standard errors are computed under the worst-case assumption that all observations within a country may be arbitrarily correlated; with 40 years per country, this worst-case adjustment inflates the estimated variance substantially relative to the empirically-relevant short-run dependence structure. The shift from the MBB to the two-way clustered standard error mechanically shrinks the implied t-statistic on $\hat{\beta}_2$ from +4.5 to +1.6. We adopt the MBB throughout the paper.

Table 6: Comparison of Standard Error Methods for the Baseline Panel Regression

This table reports the baseline panel regression coefficients (Table 5, Column 2) together with standard errors computed under four different methods: heteroskedasticity-robust, country cluster, two-way cluster (country and year), and the 2-year moving block bootstrap within country (B = 1,000 replications). Coefficients are identical across methods; only the standard errors differ. t-statistics are reported in square brackets.

	β_1 (cyclone D_{it})	β_2 ($D_{it} \times \tilde{\lambda}_{i,t-1}$)
Coefficient	-0.9798	+0.8222
Standard errors (and t-statistics in brackets):		
Heteroskedasticity-robust	(0.0326) [-30.10]	(0.2086) [+3.94]
Country cluster	(0.0880) [-11.14]	(0.5104) [+1.61]
Two-way cluster (country, year)	(0.0866) [-11.32]	(0.5023) [+1.64]
2-year moving block bootstrap	(0.0278) [-35.30]	(0.1836) [+4.48]
Observations	1,872	1,872
Countries	48	48

The remainder of this section examines whether the risk-dependent attenuation captured by $\hat{\beta}_2$ reflects genuine adaptation rather than confounding dynamics such as lagged disaster effects or trending omitted variables.

5.2 Robustness to Lagged Disasters and Country-Specific Trends

We now consider two threats to the identification of β_2 that have been highlighted in the literature on climate damages. First, past cyclones may have lagged effects on GDP growth that operate independently of adaptation behavior—reconstruction activity, capital replacement, or carry-over from previous-year output dynamics. If past disasters mechanically affect current growth and $\tilde{\lambda}_{i,t-1}$ is positively correlated with the recent disaster history of country i, then β_2 could partly reflect these lagged direct effects rather than genuine risk-dependent adaptation. Second, since the time-varying arrival rate $\hat{\lambda}_{it}$ is by construction a smooth country-specific trend, any unmodeled smooth trend in country-level GDP growth that correlates with $\hat{\lambda}_{it}$ could confound β_2. Differential country-specific growth trajectories (associated with development, demographic transitions, or other slow-moving factors) are a particular concern.

Table 7 reports panel regressions that address both threats directly. Column (1) reproduces the baseline specification from Table 5, Column (3). Columns (2) and (3) progressively add the lagged disaster indicators $D_{i,t-1}$ and $D_{i,t-2}$ as additional regressors. If β_2 were spuriously capturing the effect of recent past disasters, including these lagged D controls should substantially reduce its magnitude or significance. Column (4) instead adds country-specific linear trends $\delta_i \cdot t$, absorbing any smooth country-level evolution of GDP growth that is collinear with the trend component of $\hat{\lambda}_{it}$. Column (5) combines both extensions, including lagged disasters and country-specific trends simultaneously.

Two patterns are striking. First, $\hat{\beta}_2$ is essentially invariant across all five columns. The point estimate is +0.86 in the baseline (Column 1) and rises slightly to +0.95 in the kitchen-sink specification (Column 5); the t-statistic is between +4.9 and +5.0 in every column, and statistical significance at well below the 1% level is preserved throughout. Adding controls that explicitly capture lagged disaster dynamics or smooth country-level GDP trends therefore does not weaken the risk-dependent adaptation result; if anything, the magnitude grows modestly. Second, the lagged-disaster coefficients $L.D$ and $L^2.D$ are individually small (all less than 0.01 in absolute value) and statistically indistinguishable from zero in every specification. Past cyclones do not appear to have detectable lagged direct effects on GDP growth beyond the contemporaneous response captured by β_1. The standalone γ on $\tilde{\lambda}_{i,t-1}$ also remains statistically insignificant across all columns, confirming that climate risk does not have a direct effect on GDP growth outside of cyclone arrival years.

Taken together, these results indicate that the risk-dependent attenuation captured by β_2

Table 7: Robustness of Risk-Dependent Adaptation to Lagged Disasters and Country-Specific Trends

This table extends the baseline panel regression in Table 5, Column (3) by adding lagged cyclone indicators and country-specific linear trends. Column (1) reproduces the baseline. Columns (2) and (3) add one and two lags of the cyclone indicator, $D_{i,t-1}$ and $D_{i,t-2}$, to control for delayed direct effects of past disasters on GDP growth. Column (4) adds country-specific linear trends $\delta_i \cdot t$ to absorb any smooth country-level evolution of GDP growth that might correlate with $\tilde{\lambda}_{it}$ (which is itself a smooth trend by construction). Column (5) combines both extensions. The coefficient β_2 on $D_{it} \times \tilde{\lambda}_{i,t-1}$ remains positive and statistically significant across all specifications, indicating that the risk-dependent attenuation is not driven by lagged disaster dynamics or by country-specific GDP-growth trends. Standard errors (in parentheses) are from a 2-year moving block bootstrap within country, $B = 500$ replications. ***, **, * denote significance at the 1%, 5%, and 10% levels.

	(1)	(2)	(3)	(4)	(5)
β_1 (cyclone D_{it})	-0.9839***	-0.9839***	-0.9956***	-0.9908***	-0.9961***
	(0.0272)	(0.0272)	(0.0269)	(0.0276)	(0.0270)
γ ($\tilde{\lambda}_{i,t-1}$)	-0.0984	-0.0967	-0.0413	-0.2182*	-0.0973
	(0.0658)	(0.0693)	(0.0675)	(0.1232)	(0.1066)
β_2 ($D_{it} \times \tilde{\lambda}_{i,t-1}$)	+0.8590***	+0.8601***	+0.9415***	+0.8835***	+0.9495***
	(0.1757)	(0.1756)	(0.1878)	(0.1804)	(0.1893)
$L.D$ ($D_{i,t-1}$)		-0.0022	-0.0000		+0.0016
		(0.0154)	(0.0150)		(0.0158)
$L^2.D$ ($D_{i,t-2}$)			-0.0094		-0.0080
			(0.0152)		(0.0157)
Country FE	Yes	Yes	Yes	Yes	Yes
Region $\times$ Year FE	Yes	Yes	Yes	Yes	Yes
Country trends ($\delta_i t$)	No	No	No	Yes	Yes
Observations	1,872	1,872	1,824	1,872	1,824
Countries	48	48	48	48	48

is not driven by lagged disaster dynamics or by country-specific GDP trends. The adaptation channel is identified by within-country, within-year variation in $\hat{\lambda}_{it}$ that is orthogonal both to recent disaster history and to smooth country-level GDP evolution.

Interaction with country-specific GDP-growth trend. Barreca, Clay, Deschenes, Greenstone, and Shapiro (2016) and recent work cited in our related literature stress that confounding country-specific dynamics in the outcome—a country simultaneously growing faster and becoming more cyclone-prone, say—could mechanically produce a positive interaction between D_{it} and a smooth risk trend. To address this concern directly, Table 8 adds an interaction $D_{it} \times (\hat{g}_i \cdot t)$, where $\hat{g}_i$ is the slope of country i's real GDP per capita growth on time. This term absorbs the joint variation between cyclone exposure and any country-specific GDP-growth trajectory. The coefficient on $D_{it} \times \tilde{\lambda}_{i,t-1}$ remains +0.835 (MBB SE 0.185, $t = +4.53$, $p < 0.001$)—essentially unchanged from the baseline +0.859—and the new interaction itself is statistically zero (+0.005, $t = +0.79$). Risk-dependent adaptation is not confounded with country-specific GDP trends interacted with cyclone exposure.

Table 8: Robustness to D_{it} Interacted with Country-Specific GDP-Growth Trend

This table adds an extra interaction term $D_{it} \times (\hat{g}_i \cdot t)$ to the baseline panel regression, where $\hat{g}_i$ is the OLS slope of country i's real GDP per capita growth on time. The interaction absorbs any country-specific GDP-growth trajectory that might be confounded with cyclone exposure. Country fixed effects and region × year fixed effects are included. Standard errors (in parentheses) are 2-year moving block bootstrap within country ($B = 1000$). ***, **, * denote significance at the 1%, 5%, and 10% levels.

	(6) $+ D_{it} \times \hat{g}_i t$
β_1 (D_{it})	-0.9863^{***}
	(0.0287)
γ ($\tilde{\lambda}_{i,t-1}$)	-0.1003
	(0.0645)
β_2 ($D_{it} \times \tilde{\lambda}_{i,t-1}$)	$+0.8354^{***}$
	(0.1846)
$D_{it} \times \hat{g}_i t$	$+0.0048$
	(0.0061)
Country FE	Yes
Region × Year FE	Yes
Observations	1,872
Countries	48

Sub-sample robustness. A further concern is that the adaptation channel may be identified solely from late-sample variation, where climate trends have steepened. Table 9 re-estimates the baseline regression on two non-overlapping sub-samples: 1980–1999 and 2000–2019. The interaction $\hat{\beta}_2$ is positive and statistically significant in both: +0.706 ($t = +2.68$, $p = 0.012$) in the pre-2000 sub-sample and +1.101 ($t = +4.14$, $p < 0.001$) post-2000. The somewhat larger post-2000 magnitude is consistent with the climate-change story—as $\tilde{\lambda}_{it}$ has trended further from each country's 1980 baseline, the slope of the adaptation function is identified off a wider range of risk variation—but the channel is detectable and significant in the first half of the sample alone. The adaptation channel is not a feature of any particular sub-period.

Table 9: Sub-sample Robustness: Pre-2000 versus Post-2000

This table re-estimates the baseline panel regression on two non-overlapping sub-samples, 1980–1999 and 2000–2019, to verify that the risk-dependent adaptation channel is stable across the sample period rather than identified solely from late-sample variation. Country fixed effects and region × year fixed effects are included. Standard errors (in parentheses) are 2-year moving block bootstrap within country ($B = 1000$). ***, **, * denote significance at the 1%, 5%, and 10% levels.

	1980–1999	2000–2019
β_1 (D_{it})	-0.9694^{***}	-1.0008^{***}
	(0.0399)	(0.0386)
γ ($\tilde{\lambda}_{i,t-1}$)	-0.2447	$+0.0002$
	(0.1663)	(0.1126)
β_2 ($D_{it} \times \tilde{\lambda}_{i,t-1}$)	$+0.7065^{**}$	$+1.1011^{***}$
	(0.2634)	(0.2658)
Country FE	Yes	Yes
Region × Year FE	Yes	Yes
Observations	912	960
Countries	48	48

Local-projection test of permanent versus transitory effects. A natural question is whether the cyclone damage captured by β_1 is persistent or transitory in net growth, and whether the adaptation offset β_2 operates only at impact or lingers over subsequent years. We address both questions by running the Jordà (2005) local-projection version of the baseline regression at horizons $h = 0, 1, \ldots, 5$:

$$g^n_{i,t+h} = \mu_i^{(h)} + v_t^{(h)} + \beta_1^{(h)} D_{it} + \beta_2^{(h)} \tilde{\lambda}_{i,t-1} D_{it} + \varepsilon_{i,t+h}, \qquad h = 0, 1, \ldots, 5. \tag{15}$$

The dependent variable is net growth $g^n_{i,t+h}$ at each horizon, preserving the same units and decomposition as the baseline specification. The coefficient $\beta_1^{(h)}$ traces out whether a cyclone at t continues to reduce net growth at $t+1, t+2, \ldots$: if $\beta_1^{(h)}$ remains negative at $h \geq 1$, the damage persists in the growth rate, compounding the level loss over time; if $\beta_1^{(h)}$ reverts to zero, subsequent growth makes up the initial loss. For cyclones—which physically destroy capital—theoretical priors favor persistence, since lost capital stock does not rebuild instantaneously. The coefficient $\beta_2^{(h)}$ traces out the multi-year dynamics of the adaptation channel: whether the risk-dependent offset is concentrated at impact or extends to the recovery path. Country fixed effects, region $\times$ year fixed effects, and 2-year MBB standard errors are unchanged from the baseline at each horizon.

Figure 3 reports the results, with detailed numbers in Table 10. The left panel shows the direct cyclone effect $\hat{\beta}_1^{(h)}$ on net growth at each horizon. At impact ($h = 0$), $\hat{\beta}_1^{(0)} = -0.984$, identical to the baseline by construction. At every subsequent horizon $h = 1, \ldots, 5$, the point estimates are essentially zero and statistically insignificant. Crucially, there is no evidence of above-normal growth at any post-disaster horizon—the rebound that would characterize a transitory level effect is absent. The economy absorbs the growth loss at impact and then resumes growing at its pre-disaster rate from a permanently lower level. This is the signature of a permanent level effect: the output lost to cyclone damage is never recovered through subsequent catch-up growth. The result is consistent with the physical destruction of capital stock, which depresses the level of output permanently unless replaced by new investment—and since we measure g^n net of the investment-depreciation margin, the rebuilding channel is already stripped out. The right panel shows the risk-dependent adaptation channel $\hat{\beta}_2^{(h)}$: at impact ($h = 0$), $\hat{\beta}_2^{(0)} = +0.859$ ($t = +4.89$, $p < 0.001$)—identical to the baseline by construction. At every future horizon $h \geq 1$, $\hat{\beta}_2^{(h)}$ is statistically indistinguishable from zero ($|\hat{\beta}_2^{(h)}| \leq 0.28$ in absolute value, with $|t| \leq 1.12$). The adaptation channel operates contemporaneously: a cyclone in year t produces its attenuated damage in year t, and there is no detectable lingering effect on g^n in years $t+1, t+2, \ldots, t+5$. This is consistent with adaptation working through pre-positioned protective capital—seawalls, building codes, early warning systems—that reduces damage when the cyclone strikes, rather than through post-disaster recovery dynamics.

5.3 Placebo Tests

The robustness checks in Section 5.2 demonstrate that $\hat{\beta}_2$ does not move under additional parametric controls. We now complement those tests with a non-parametric, design-based identification check: placebo permutations that disrupt the link between cyclone arrivals, prior

Table 10: Local-Projection Estimates of Cyclone Damage at Horizons $h = 0, 1, \ldots, 5$

This table reports local-projection (Jordà, 2005) estimates at horizons $h = 0, 1, \ldots, 5$. Each row estimates $g^n_{i,t+h} = \beta_1^{(h)} D_{it} + \beta_2^{(h)} D_{it} \times \tilde{\lambda}_{i,t-1} + \tilde{\lambda}_{i,t-1} + \alpha_i + \nu_{r(i),t} + \varepsilon_{i,t+h}$. Column $\hat{\beta}_1^{(h)}$ traces the persistence of the direct cyclone damage in net growth: negative values at $h \geq 1$ indicate that cyclone damage compounds in the growth rate rather than reverting. Column $\hat{\beta}_2^{(h)}$ traces the persistence of the risk-dependent adaptation channel. All specifications include country fixed effects, region × year fixed effects, and the standalone $\tilde{\lambda}_{i,t-1}$ on the RHS (coefficients not shown). Standard errors (in parentheses) are 2-year moving block bootstrap within country ($B = 500$). ***, **, * denote significance at the 1%, 5%, and 10% levels.

	Direct cyclone damage		Risk-dependent adaptation		
	$\hat{\beta}_1^{(h)}$ (D_{it})	(SE)	$\hat{\beta}_2^{(h)}$ $(D_{it} \times \tilde{\lambda}_{i,t-1})$	(SE)	N
$h = 0$	-0.984^{***}	(0.027)	$+0.859^{***}$	(0.176)	1,872
$h = 1$	$+0.028$	(0.032)	-0.035	(0.257)	1,824
$h = 2$	$+0.052$	(0.036)	$+0.104$	(0.236)	1,776
$h = 3$	$+0.053^{*}$	(0.032)	-0.043	(0.262)	1,728
$h = 4$	$+0.029$	(0.034)	$+0.278$	(0.248)	1,680
$h = 5$	$+0.032$	(0.040)	-0.009	(0.271)	1,632

Figure 3: Local-projection impulse responses at horizons $h = 0, 1, \ldots, 5$

This figure plots local-projection estimates of the impact of a cyclone on net GDP growth $g^n_{i,t+h}$ at horizons $h = 0, 1, \ldots, 5$ years out. The left panel shows the direct cyclone coefficient $\hat{\beta}_1^{(h)}$; the right panel shows the risk-dependent adaptation coefficient $\hat{\beta}_2^{(h)}$ on $D_{it} \times \tilde{\lambda}_{i,t-1}$. Each estimate comes from a separate regression with country fixed effects, region × year fixed effects, and 2-year moving-block bootstrap standard errors. Vertical bars are 95% confidence intervals.

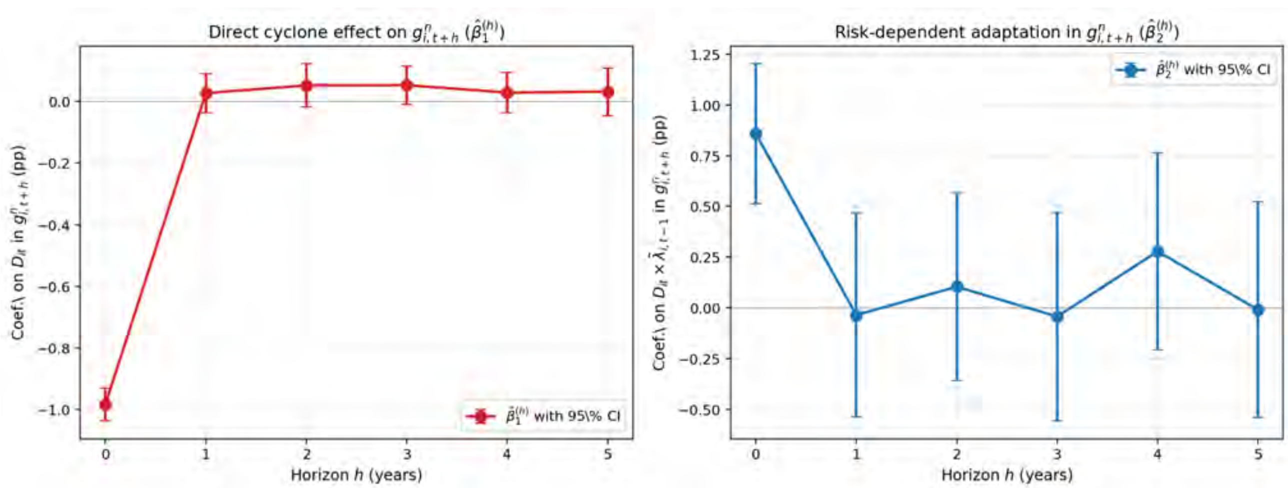

risk, and country growth dynamics. If our baseline result reflects a genuine country-specific relationship between D_{it}, $\tilde{\lambda}_{i,t-1}$, and g^n_{it}, then breaking that link should drive the estimated $\hat{\beta}_2$ toward zero. If instead the result reflects a generic feature of cyclone-prone country panels—for example, a mechanical interaction between any disaster sequence and any smooth trend—then placebo $\hat{\beta}_2$ should remain in the neighborhood of the actual estimate.

We construct two placebos. In Panel A (Placebo Disasters), for each replication we draw a random derangement of the 58 countries and reassign each country's full cyclone time series $\{D_{it}\}$ to a different country. Each country's own $\hat{\lambda}_{i,t-1}$ and g^n_{it} are preserved. We then re-estimate the baseline panel regression (Table 5 Column (3)) on the resulting placebo dataset. The procedure breaks the link between a country's own disaster history and its growth dynamics while leaving the country's own risk trajectory intact. In Panel B (Placebo Risk), we instead permute the $\hat{\lambda}_{i,t-1}$ series across countries while preserving each country's own D_{it} and g^n_{it}. This breaks the link between a country's own risk evolution and the disaster-conditional damage function. We run $N = 1{,}000$ permutations in each panel.

Figure 4: Placebo Distributions of $\hat{\beta}_2$

Distributions of the adaptation coefficient $\hat{\beta}_2$ under two placebo schemes, each with $N = 1{,}000$ random derangements of country assignments. Panel A reassigns each country's cyclone series $\{D_{it}\}$ to another country while preserving the original country's $\hat{\lambda}_{i,t-1}$ and g^n_{it}. Panel B reassigns each country's $\hat{\lambda}_{i,t-1}$ series. The vertical red line marks the actual $\hat{\beta}_2 = +0.54$ from Table 5 Column (3). The right-tail p-values are the fraction of placebo estimates at least as large as the actual.

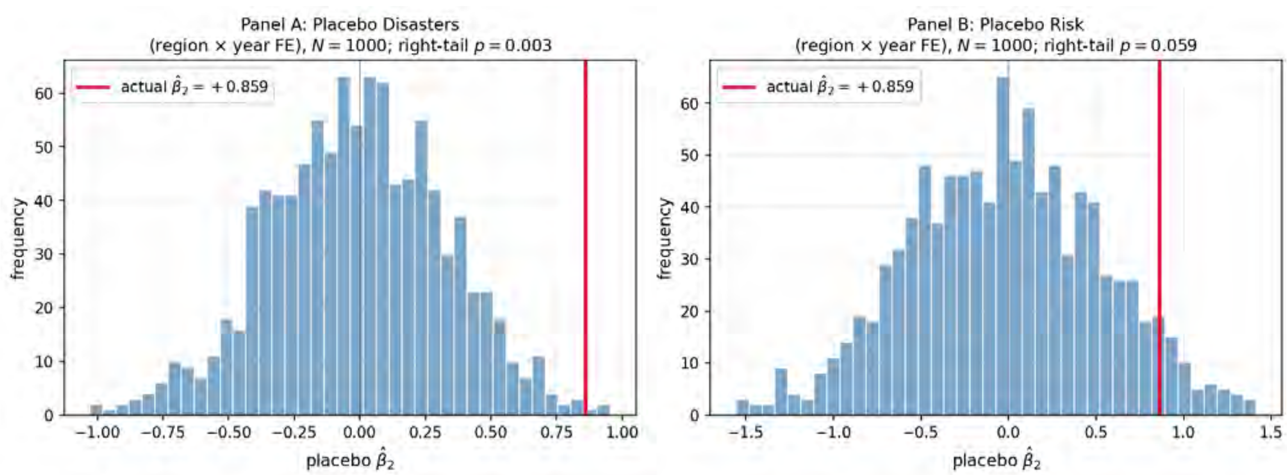

Figure 4 and Table 11 report the placebo distributions. Two patterns are notable. First, both placebo distributions are centered at zero (Panel A mean -0.009; Panel B mean -0.024), exactly as expected under the null that there is no systematic relationship between the permuted

Table 11: Placebo Tests: Distribution of $\hat{\beta}_2$ under Random Reassignment

This table summarizes the placebo distribution of the adaptation coefficient $\hat{\beta}_2$ under two randomization schemes, using the baseline specification with country and region $\times$ year fixed effects. Panel A (Placebo Disasters) randomly permutes the cyclone time series across countries (using a derangement, so no country retains its own series), holds each country's own $\hat{\lambda}_{i,t-1}$ and growth fixed, and re-estimates the baseline panel regression. Panel B (Placebo Risk) does the analogous permutation of the $\hat{\lambda}_{i,t-1}$ series across countries while preserving each country's own disaster sequence. Each panel reports the mean and standard deviation of the placebo distribution, the actual $\hat{\beta}_2$ from Table 5 Column (3), and the right-tail permutation p-value (fraction of placebo $\hat{\beta}_2^*$ at least as large as the actual). $N = 1{,}000$ permutations per panel.

	Placebo Disasters	Placebo Risk
Mean of placebo $\hat{\beta}_2^*$	-0.0091	-0.0235
Standard deviation of placebo $\hat{\beta}_2^*$	0.3339	0.5508
Actual $\hat{\beta}_2$ (Table 5, Column 3)	+0.8590	+0.8590
Right-tail p-value (one-sided)	0.0030^{***}	0.0590^{*}
Two-sided p-value	0.0090^{***}	0.1220
Percentile of actual	99.7%	94.1%
N permutations	1,000	1,000

variables and growth dynamics. This confirms that the baseline panel structure—country and region $\times$ year fixed effects, the disaster indicator, and the standalone $\tilde{\lambda}$—is well behaved and does not mechanically produce a positive $\hat{\beta}_2$. Second, the actual $\hat{\beta}_2 = +0.86$ lies in the upper tail of both placebo distributions. Under Placebo Disasters (Panel A), the actual estimate is at the 99.7th percentile, with a right-tail p-value of 0.003. The result is therefore essentially impossible to reproduce by random reassignment of cyclone series: only 3 of 1,000 placebo permutations yield a $\hat{\beta}_2$ as large as the actual estimate. The country's own disaster sequence is essential.

Under Placebo Risk (Panel B), the actual estimate is at the 94th percentile with a right-tail p-value of 0.059, significant at the conventional 10% level. The Placebo Risk distribution is wider than Placebo Disasters (SD = 0.55 vs 0.33), reflecting that any country-specific trend in $\tilde{\lambda}_{i,t-1}$, when interacted with a country's actual disasters, retains some trend-times-disaster variation even after the region $\times$ year fixed effects absorb the shared regional component. Even under this demanding null, the actual estimate is clearly in the right tail. Taken together, the two placebos indicate that both the country's own cyclone realizations and the country's specific risk trajectory contribute to the identification of $\hat{\beta}_2$: random reassignment of either component substantially weakens the estimated adaptation effect, and random reassignment of disasters effectively eliminates it.

Non-cyclone-prone country placebo. A complementary placebo design takes the cyclone-affected sample out of the picture entirely. For each replication we randomly assign each non-cyclone-prone country—a country with zero landfalls in 1980–2019—a cyclone-affected country's full sequences of D_{it} and $\tilde{\lambda}_{i,t-1}$ as "borrowed" cyclone exposure, then re-estimate the baseline panel regression on the non-cyclone country's own growth data. Because the assignment of cyclone exposure to non-cyclone countries is random and the growth data are real, any non-zero placebo $\hat{\beta}_2^*$ must come from a generic mechanical interaction between cyclone-shaped series and growth panels rather than from any country-specific adaptation response. Across $N = 500$ valid permutations of the 50 non-cyclone-prone countries, the placebo distribution of $\hat{\beta}_2^*$ is essentially centered at zero (mean -0.004, SD 0.291), and the actual $\hat{\beta}_2 = +0.859$ lies at the 100th percentile with a right-tail p-value below 0.0001 (Table 12). The placebo and the actual share no overlap in the relevant range. This is the sharpest design-based check available: it relies on neither the cyclone-affected sample nor the parametric MBB inference, and it decisively rejects the hypothesis that the baseline result is a generic feature of cyclone-shaped regressors interacted with growth panels.

Table 12: Non-Cyclone-Prone Country Placebo Test

This table reports a sharp placebo test that disrupts the link between cyclone arrivals and growth dynamics at the country level. For each placebo replication, we randomly assign each non-cyclone-prone country (a country with zero landfalls in 1980–2019) a cyclone-affected country's full sequences of D_{it} and $\tilde{\lambda}_{i,t-1}$ as "borrowed" cyclone exposure, and re-estimate the baseline panel regression on the non-cyclone countries' own growth data. Because the assignment of cyclone exposure to non-cyclone countries is random and the growth data are real, any positive placebo $\hat{\beta}_2^*$ must come from a generic mechanical interaction between cyclone-shaped series and growth panels rather than from a country-specific relationship. We run $N = 500$ valid permutations. ***, **, * denote significance at the 1%, 5%, and 10% levels.

	Non-cyclone-prone placebo
Mean of placebo $\hat{\beta}_2^*$	-0.0043
Standard deviation of placebo $\hat{\beta}_2^*$	0.2909
Actual $\hat{\beta}_2$ (Table 5 Column 3)	$+0.8590$
Right-tail p-value (one-sided)	$< 0.0001^{***}$
Percentile of actual	100.0%
Non-cyclone-prone countries	50
Country-years available	$1,950$
Placebo permutations	500

5.4 Horse Race against the Cumulative-Exposure Moderator

A natural concern is that our time-varying NHPP risk path $\tilde{\lambda}_{i,t-1}$ is, in practice, a relabeled version of the time-invariant cumulative-exposure moderator n_i that the prior adaptation literature uses—most prominently Hsiang and Jina (2014) and Bakkensen and Mendelsohn (2016), who augment the basic disaster regression with $D_{it} \times n_i$, where n_i is country i's historical cyclone frequency. If the cross-country contrast in n_i already absorbs the identifying variation, then β_2 in our specification offers little beyond a different parameterization of the same fact.

We address this concern with a direct horse race. Define $n_i = (1/T_i)\sum_t D_{it}$ as country i's mean cyclone-landfall indicator over 1980–2019—the time-invariant cumulative-exposure moderator in the spirit of the prior literature. We estimate three nested specifications, all with country fixed effects, region × year fixed effects, and 2-year MBB standard errors:

1. HJ-style only: $g^n_{it} = \alpha_i + \nu_{r,t} + \beta_1 D_{it} + \psi(D_{it} \times n_i) + \varepsilon_{it}$.
2. Our spec only: $g^n_{it} = \alpha_i + \nu_{r,t} + \beta_1 D_{it} + \gamma\tilde{\lambda}_{i,t-1} + \beta_2(D_{it} \times \tilde{\lambda}_{i,t-1}) + \varepsilon_{it}$ (= Table 5 Column 3).
3. Both moderators jointly: the full specification with both $D_{it} \times n_i$ and $D_{it} \times \tilde{\lambda}_{i,t-1}$ interactions, plus the standalone level of $\tilde{\lambda}$.

Table 13 reports the results. Two findings stand out.

First, the HJ-style moderator ψ on $D_{it} \times n_i$ is statistically insignificant on its own in Column (1): $\hat{\psi} = +0.133$ with MBB standard error 0.118 and $p = 0.26$. The pure cross-country cumulative-exposure contrast does not survive the modern identification controls in our sample: with region × year fixed effects absorbing shared regional trends and the panel restricted to the 48 cyclone-affected countries with region assignment, n_i alone does not pick up a statistically reliable adaptation pattern. This is consistent with the broader concern in the recent identification literature that early adaptation findings relied on relatively loose fixed-effect structures.

Second, and more important: $\hat{\beta}_2$ on $D_{it} \times \tilde{\lambda}_{i,t-1}$ survives the horse race essentially unchanged. Adding $D_{it} \times n_i$ to the regression moves the point estimate of β_2 from +0.859 (Column 2) to +0.832 (Column 3)—a 3% change, well within its MBB standard error—while ψ in the joint specification falls further to +0.077 with $p = 0.54$. Conditional on our time-varying NHPP risk moderator, the cross-country cumulative-exposure contrast adds nothing; conditional on n_i, our time-varying moderator retains essentially all of its explanatory power.

The interpretation is straightforward. Our moderator $\tilde{\lambda}_{i,t-1} = \hat{\lambda}_{i,t-1} - \hat{\lambda}_{i,0}$ is by construction the deviation of current risk from the country's baseline level. A country that has always faced high cyclone exposure (high n_i) but with a stable risk trajectory has $\tilde{\lambda} \approx 0$, and the model predicts no incremental adaptation pressure beyond its baseline investment. A country whose risk has been rising has $\tilde{\lambda} > 0$, and the adaptation response is identified off precisely this incremental risk. The horse race shows that the empirical content of β_2 is "adaptation responds to changes in risk," not the static "high-exposure countries adapt more" fact that n_i would deliver. This is the empirical hook that motivates Section 6's forward-looking projections: if climate change pushes $\hat{\lambda}_{it}$ above its country's historical baseline, the adaptation gain accruing to that country is governed by β_2 on the rising $\tilde{\lambda}$ path, not by its fixed historical exposure.

Table 13: Horse Race: Time-Invariant Cumulative-Exposure Moderator n_i vs Time-Varying NHPP Risk $\tilde{\lambda}_{i,t-1}$

This table reports a horse race between the time-invariant Hsiang–Jina-style moderator n_i (country i's mean cyclone landfall indicator over 1980–2019) and our time-varying NHPP risk moderator $\tilde{\lambda}_{i,t-1}$. Column (1) is the HJ-style specification: $g^n_{it} = \alpha_i + \nu_{r,t} + \beta_1 D_{it} + \psi(D_{it} \times n_i) + \varepsilon_{it}$. Column (2) is our baseline specification with $\tilde{\lambda}_{i,t-1}$ as the moderator. Column (3) includes both interactions jointly. Country fixed effects and region × year fixed effects are included throughout. Standard errors (in parentheses) are 2-year moving block bootstrap within country, $B = 1000$. The horse race tests whether the time-varying NHPP risk path delivers identifying variation beyond the cross-country cumulative-exposure contrast that the older literature exploits. ***, **, * denote significance at the 1%, 5%, and 10% levels.

	(1) HJ-style	(2) Our spec	(3) Horse race
β_1 (cyclone D_{it})	-0.9753^{***}	-0.9839^{***}	-1.0067^{***}
	(0.0446)	(0.0279)	(0.0448)
ψ (HJ: $D_{it} \times n_i$)	+0.1333		+0.0766
	(0.1176)		(0.1219)
γ ($\tilde{\lambda}_{i,t-1}$)		−0.0984	−0.0976
		(0.0643)	(0.0645)
β_2 ($D_{it} \times \tilde{\lambda}_{i,t-1}$)		$+0.8590^{***}$	$+0.8319^{***}$
		(0.1839)	(0.1905)
Country FE	Yes	Yes	Yes
Region × Year FE	Yes	Yes	Yes
Observations	1,872	1,872	1,872
Countries	48	48	48

5.5 TWFE Robustness: A Binned Dose-Response Estimator

The baseline specification in Equation (14) is a two-way fixed-effects regression with a continuous moderator $\tilde{\lambda}_{i,t-1}$. A recent literature has shown that the TWFE coefficient on a continuous treatment-intensity term is generally not the average causal response (ACR), but a weighted average across units with potentially negative or oddly-concentrated weights when treatment effects are heterogeneous (de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021; Borusyak, Jaravel, and Spiess, 2024; Callaway, Goodman-Bacon, and Sant'Anna, 2024). To verify that our $\hat{\beta}_2$ is not an artifact of this weighting, we implement a non-parametric binned dose-response estimator in the spirit of Callaway, Goodman-Bacon, and Sant'Anna (2024)'s ACR for continuous treatments.

The procedure has three steps. First, we sort all cyclone-affected country-years by the lagged risk $\tilde{\lambda}_{i,t-1}$ and assign them to $K = 8$ equal-frequency quantile bins. Second, within each bin k we estimate the cyclone disaster ATT as the coefficient on D_{it} in a within-country regression restricted to that bin's observations:

$$g_{it}^n = \alpha_i + \tau_k D_{it} + \varepsilon_{it}, \qquad (i,t) \in \text{bin } k, \tag{16}$$

with country-clustered standard errors. Each $\hat{\tau}_k$ is independently identified from a clean within-country comparison among years in bin k's risk range; the bin-level estimator avoids the cross-bin weighting that drives the TWFE critique. Third, we recover the dose-response slope by a weighted least squares regression of the bin-level ATTs $\hat{\tau}_k$ on the bin-mean $\bar{\tilde{\lambda}}_k$, using weights $1/\widehat{\mathrm{Var}}(\hat{\tau}_k)$. The slope is a non-parametric estimate of β_2 and is the discrete-bin analog of the average causal response slope in Callaway, Goodman-Bacon, and Sant'Anna (2024); inference on the binned scatter follows the framework of Cattaneo, Crump, Farrell, and Feng (2024).

Table 14 reports the eight bin-level ATTs and the dose-response slope. The disaster ATT smoothly improves with the lagged risk level, from $\hat{\tau}_1 = -1.060$ percentage points in the lowest-risk bin to $\hat{\tau}_8 = -0.797$ percentage points in the highest-risk bin—countries facing higher cyclone risk experience smaller GDP-growth damages from a given landfall, exactly the risk-dependent adaptation channel that β_2 is designed to identify. The implied slope of $\hat{\tau}_k$ on the bin mean of $\tilde{\lambda}$ is +0.683 (SE 0.249, $t = +2.74$, $p < 0.01$), in the same direction and similar magnitude as the panel point estimate of +0.859 from Table 5. Figure 5 plots the bin-level ATTs against the bin means together with the WLS fit; the relationship is visibly linear, validating both the panel specification's functional form and the conclusion that the panel $\hat{\beta}_2$ captures a real cross-country dose-response rather than a TWFE weighting artifact.

Table 14: TWFE-Robust Binned Dose-Response Estimator of Risk-Dependent Adaptation

This table reports a TWFE-robust binned dose-response estimator of the risk-dependent adaptation coefficient β_2 on $D_{it} \times \tilde{\lambda}_{i,t-1}$. Country-years are sorted by the lagged risk variable $\tilde{\lambda}_{i,t-1}$ into 8 quantile bins. Within each bin k, we estimate the cyclone disaster ATT $\hat{\tau}_k$ as the coefficient on D_{it} in a within-country regression $g^n_{it} = \alpha_i + \tau_k D_{it} + \varepsilon_{it}$ restricted to the bin's observations, with country-clustered standard errors. Each $\hat{\tau}_k$ is therefore identified from a clean within-country comparison among years in the bin's risk range and does not depend on the TWFE weighting that is the focus of the recent critique (de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021; Borusyak, Jaravel, and Spiess, 2024; Callaway, Goodman-Bacon, and Sant'Anna, 2024). The dose-response slope at the bottom of the table is the WLS coefficient of $\hat{\tau}_k$ on the bin-mean $\tilde{\lambda}$ (weights $1/\text{SE}(\hat{\tau}_k)^2$); it is the discrete-bin analog of the average causal response slope in Callaway, Goodman-Bacon, and Sant'Anna (2024) and provides a non-parametric estimate of β_2. Compare to the panel baseline $\hat{\beta}_2 = +0.859$ (Table 5, Column 3). ***, **, * denote significance at the 1%, 5%, and 10% levels.

Bin	Range of $\tilde{\lambda}_{i,t-1}$	Mean $\tilde{\lambda}$	Obs.	Disasters	$\hat{\tau}_k$	(SE)	
$k = 1$	$[-0.153, -0.044]$	-0.068	234	40	-1.060^{***}	(0.122)	
$k = 2$	$[-0.044, -0.019]$	-0.033	234	28	-0.963^{***}	(0.165)	
$k = 3$	$[-0.018, -0.004]$	-0.010	234	53	-1.073^{***}	(0.201)	
$k = 4$	$[-0.004, +0.005]$	$+0.001$	234	49	-0.900^{***}	(0.108)	
$k = 5$	$[+0.005, +0.031]$	$+0.013$	234	64	-1.044^{***}	(0.171)	
$k = 6$	$[+0.031, +0.080]$	$+0.062$	234	152	-0.876^{***}	(0.222)	
$k = 7$	$[+0.081, +0.191]$	$+0.123$	234	105	-0.953^{***}	(0.164)	
$k = 8$	$[+0.192, +0.372]$	$+0.253$	234	48	-0.797^{***}	(0.156)	
Dose-response slope (WLS of $\hat{\tau}_k$ on bin-mean $\tilde{\lambda}$):					$+0.683^{***}$	(0.249)	$t = +2.74$
Panel baseline $\hat{\beta}_2$ for comparison:					$+0.859^{***}$	(0.184)	$t = +4.67$

Figure 5: Binned dose-response of cyclone damage in lagged risk $\tilde{\lambda}_{i,t-1}$

This figure plots the bin-level disaster ATT $\hat{\tau}_k$ against the bin-mean lagged risk $\bar{\tilde{\lambda}}_k$ for the $K = 8$ quantile bins in Table 14. Within each bin we estimate $\hat{\tau}_k$ as the coefficient on D_{it} in a country-fixed-effects regression of g^n_{it} on D_{it} restricted to that bin's observations; vertical bars are 95% confidence intervals based on country-clustered standard errors. The red line is the WLS fit of $\hat{\tau}_k$ on $\bar{\tilde{\lambda}}_k$ (weights $1/\text{SE}(\hat{\tau}_k)^2$); its slope +0.683 is the non-parametric TWFE-robust analog of the panel $\hat{\beta}_2$, in the spirit of the average causal response of Callaway, Goodman-Bacon, and Sant'Anna (2024).

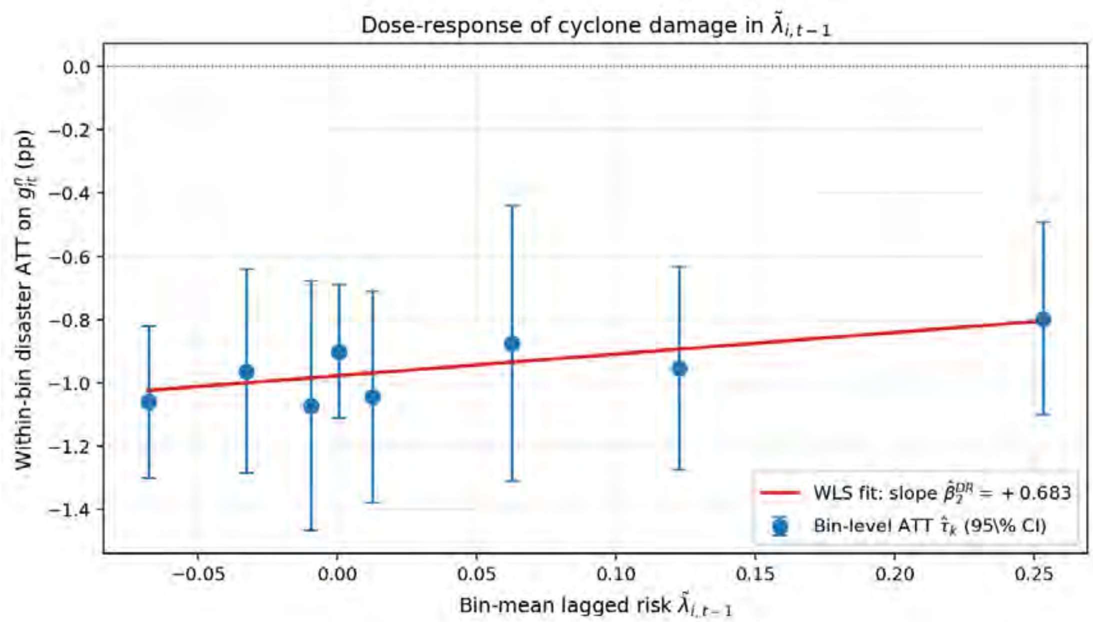

6 Counterfactuals

6.1 In Sample: 1980-2019

To quantify the cumulative welfare value of risk-dependent adaptation in our sample, we compare each country's actual 2019 real GDP per capita to the counterfactual 2019 income that would have obtained had the risk-dependent adaptation channel been absent. The local-projection estimates in Table 10 justify a simple counterfactual based on the impact-year coefficients alone: since cyclone damage to net growth is concentrated entirely at $h = 0$ with no additional effects at subsequent horizons, the counterfactual need only adjust growth in disaster years and compound forward. The counterfactual takes the actual realized growth path and strips out the estimated adaptation offset on disaster years, propagated forward from the 1980 initial condition.

Formally, for each country i we initialize $GDP^{CF}_{i,1980} = GDP_{i,1980}$ (the actual 1980 real GDP per capita) and iterate:

$$GDP^{CF}_{it} \;=\; GDP^{CF}_{i,t-1} \times \left(1 \;+\; \frac{g^{\%}_{it} \;-\; \hat{\beta}_2\,\tilde{\lambda}_{i,t-1}\,D_{it}}{100}\right), \tag{17}$$

where $g^{\%}_{it} = 100 \cdot (GDP_{it}/GDP_{i,t-1} - 1)$ is each country's actual real GDP per capita growth in percentage points, D_{it} equals one in cyclone landfall years, $\tilde{\lambda}_{i,t-1} = \hat{\lambda}_{i,t-1} - \hat{\lambda}_{i,0}$ is the deviation of the NHPP quadratic time-trend arrival rate from its 1980 baseline, and $\hat{\beta}_2 = +0.859$ (in percentage points per unit of $\tilde{\lambda}$) is the single panel coefficient estimated in Table 5 Column (3). Inside the parentheses, both terms are in percentage points; division by 100 converts to a fraction so the bracketed factor is a gross growth rate. The adaptation offset $\hat{\beta}_2\tilde{\lambda}_{i,t-1}D_{it}$ captures, in percentage points, the additional growth a country obtains in disaster years for having let its risk perceptions, and its protective stock, evolve with cumulative cyclone experience. Subtracting this term yields the no-adaptation growth path.

First, identification of the adaptation channel does not rely on country-specific coefficients, which were poorly identified given the short cyclone series available per country. We use a single, panel-wide $\hat{\beta}_2$ that is estimated under demanding controls (country FE, region × year FE), addresses serial dependence through a 2-year moving block bootstrap, and is robust to placebo permutations. Second, the time-varying risk $\hat{\lambda}_{i,t-1}$ comes from the parsimonious NHPP quadratic time-trend. Together these choices produce welfare estimates that contribute to the prior literature.

Table 15 reports the summary statistics across the 48 cyclone-affected countries with complete 1980-2019 coverage. Without risk-dependent adaptation, mean 2019 real GDP per capita would be roughly 0.5% lower, with substantial heterogeneity across countries: at the 75th percentile, the cumulative loss is about 1.1%. The countries with the largest counterfactual losses are precisely those with the highest cumulative cyclone exposure and trend-rising risk—the United States (−2.9%), South Korea (−2.4%), India (−2.1%), China (−2.0%), Costa Rica, Dominican Republic, Panama, Thailand, and Bangladesh (all between −1.2% and −2.0%). Countries with flat or declining trend risk have negligible or even slightly positive counterfactual gaps, since $\tilde{\lambda}_{i,t-1}$ is near zero or negative for these countries. When $\tilde{\lambda}_{i,t-1} < 0$—i.e., a country's estimated cyclone risk has fallen below its 1980 baseline—the adaptation offset $\hat{\beta}_2 \tilde{\lambda}_{i,t-1} D_{it}$ is negative: the model implies that the country has let its adaptive capital depreciate as risk has receded, so that when a cyclone does strike, the per-disaster damage is somewhat larger than the constant-damage benchmark. Stripping out this channel in the counterfactual therefore raises growth slightly, producing a counterfactual income that exceeds the actual. This accounts for the small negative difference at the 25th percentile in Table 15. Figure 6 visualizes the actual and counterfactual time-series paths side-by-side for all 48 countries.

The magnitudes capture only the marginal risk-dependent channel on top of the baseline cyclone effect $\hat{\beta}_1 \approx -1.0$. We interpret the result as a lower bound on the welfare value of adaptation: the panel coefficient averages across countries with very different exposure profiles, and the NHPP quadratic deviation $\tilde{\lambda}$ is a deliberately conservative measure of how much risk has changed within country during the sample.

Table 15: Summary statistics of counterfactual real GDP per capita in 2019

This table presents the summary statistics of the counterfactual real GDP per capita in 2019 across the 48 cyclone-affected countries, shutting down the risk-dependent adaptation channel. The counterfactual evolves recursively as $\text{GDP}^{CF}_{it} = \text{GDP}^{CF}_{i,t-1} \times \left(1 + (g^{\%}_{it} - \hat{\beta}_2 \tilde{\lambda}_{i,t-1} D_{it})/100\right)$, where $g^{\%}_{it} = 100 \cdot (\text{GDP}_{it}/\text{GDP}_{i,t-1} - 1)$ is the actual real GDP per capita growth in percentage points, $\hat{\beta}_2 = +0.859$ (in percentage points per unit of $\tilde{\lambda}$) is from Table 5 Column (3), and $\tilde{\lambda}_{i,t-1} = \hat{\lambda}_{i,t-1} - \hat{\lambda}_{i,0}$ is the NHPP quadratic time-trend deviation from baseline. Both terms in the numerator are in percentage points; division by 100 converts to a fraction. The recursion starts at the actual 1980 real GDP per capita. Real GDP per capita is in thousands of constant 2015 USD.

	Actual 2019	Counterfactual 2019	Difference (%)
Mean	15.63	15.55	0.50
Standard deviation	21.03	20.94	0.87
Median	5.87	5.86	0.01
25th percentile	1.89	1.86	-0.06
75th percentile	17.09	16.87	1.06
Countries	48	48	48

Figure 6: In-sample counterfactual real GDP per capita absent risk-dependent adaptation

This figure shows the in-sample counterfactual real GDP per capita compared to the actual time-series of real GDP per capita from 1980 to 2019 for the 48 cyclone-affected countries with full coverage. The left panel shows the actual paths of real GDP per capita; the right panel shows the counterfactual income paths constructed from Equation (17) with the risk-dependent adaptation channel switched off. Real GDP per capita is in thousands of constant 2015 USD.

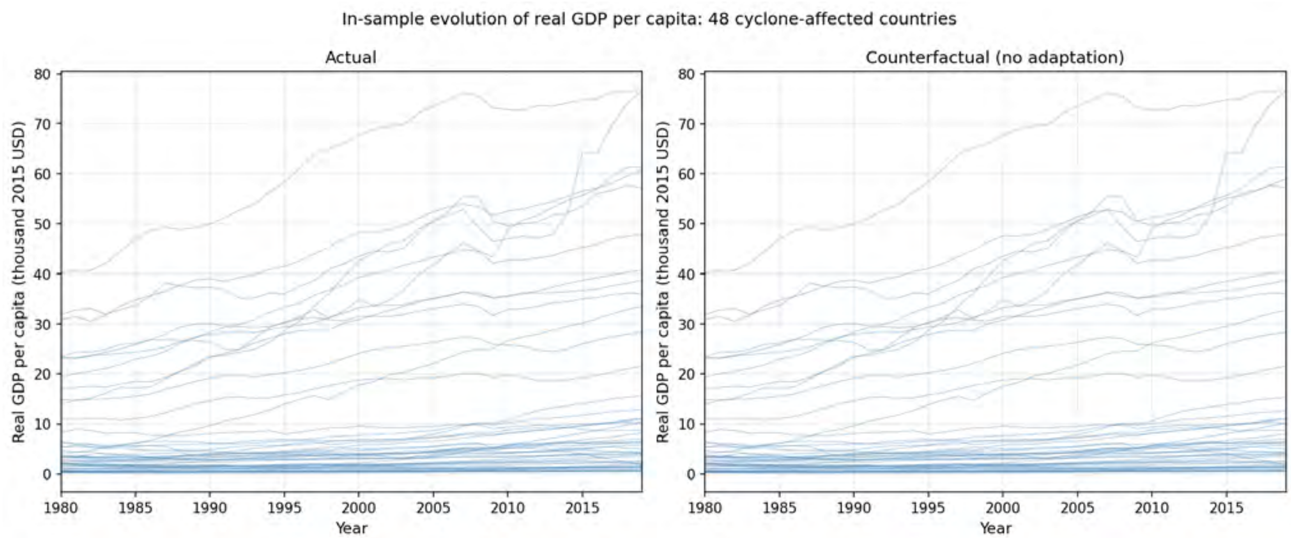

6.2 Out of Sample: 2020 to 2100

The in-sample counterfactual quantifies the cumulative gain that risk-dependent adaptation has already delivered. The more consequential exercise for policy is forward-looking: given that climate change is projected to push cyclone arrival rates above their historical baselines for many countries, how much of the resulting damage will be offset by adaptation? We answer this question by combining our estimated NHPP risk paths and the panel-baseline interaction coefficient with the standard SSP5 (Shared Socio-Economic Pathway 5, "Fossil-Fueled Development"; Kriegler, Bauer, Popp, Humpenöder, Leimbach, Strefler, Baumstark, Bodirsky, Hilaire, Klein, et al., 2017) growth projections that the climate-economy literature uses for long-horizon forecasting.

For each cyclone-affected country and each future year $t \in \{2020, \ldots, 2100\}$, the projected NHPP arrival rate $\hat{\lambda}_{it}$ is obtained by extrapolating the quadratic time-trend model of Section 4 forward.[11] Each country's projected baseline growth rate η_{it} is the SSP5 forecast in country-year

[11] The arrival-rate path is censored at 1 to respect the Bernoulli upper bound of the binary landfall indicator.

t. We then project real GDP per capita forward under three scenarios that progressively layer on climate damages and adaptation:

(A) **SSP5 baseline** (no climate damages): $GDP_{it}^{A} = GDP_{i,t-1}^{A} \times (1 + \eta_{it}/100)$.

(B) **Fixed adaptation**: cyclones inflict their baseline per-disaster damage $\hat{\beta}_1$ in every disaster year, with the expected damage in year t equal to $\hat{\beta}_1 \hat{\lambda}_{it}$: $GDP_{it}^{B} = GDP_{i,t-1}^{B} \times \big(1 + (\eta_{it} + \hat{\beta}_1 \hat{\lambda}_{it})/100\big)$.

(C) **Risk-dependent adaptation**: the per-disaster damage falls with the deviation of risk from the country's 1980 baseline, so the expected annual damage is $(\hat{\beta}_1 + \hat{\beta}_2 \tilde{\lambda}_{it})\hat{\lambda}_{it}$ where $\tilde{\lambda}_{it} = \hat{\lambda}_{it} - \hat{\lambda}_{i,0}$: $GDP_{it}^{C} = GDP_{i,t-1}^{C} \times \big(1 + (\eta_{it} + (\hat{\beta}_1 + \hat{\beta}_2 \tilde{\lambda}_{it})\hat{\lambda}_{it})/100\big)$.

The interaction coefficient $\hat{\beta}_2 = +0.859$ and the baseline cyclone effect $\hat{\beta}_1 = -0.984$ are both from Table 5, Column (3). All three paths are initialized at each country's actual 2019 real GDP per capita. The comparison of (C) to (B) isolates the value of risk-dependent adaptation: it answers the policy question of how much of the cyclone damage in a climate-stressed future will be offset by the same risk-dependent adaptation channel we have estimated from historical data.

Table 16: Projected real GDP per capita in 2050 and 2100: SSP5 + cyclone damages with and without risk-dependent adaptation

This table reports summary statistics of the projected real GDP per capita across the cyclone-affected countries in 2050 and 2100 under three scenarios. "SSP5 baseline" applies the SSP5 growth path (Kriegler, Bauer, Popp, Humpenöder, Leimbach, Strefler, Baumstark, Bodirsky, Hilaire, Klein, et al., 2017) with no climate damages. "Fixed adaptation" adds expected cyclone damages $\hat{\beta}_1\hat{\lambda}_{it}$ at each future year, holding the per-cyclone damage at the constant level $\hat{\beta}_1 = -0.984$. "Risk-dependent adaptation" adds expected damages $(\hat{\beta}_1 + \hat{\beta}_2\tilde{\lambda}_{it})\hat{\lambda}_{it}$, allowing per-cyclone damage to be smaller when risk has risen above the country's 1980 baseline ($\hat{\beta}_2 = 0.859$, both from Table 5, Column 3). $\hat{\lambda}_{it}$ is the projected NHPP arrival rate from the quadratic time-trend model (Section 4); $\tilde{\lambda}_{it} = \hat{\lambda}_{it} - \hat{\lambda}_{i,0}$. All three paths are initialized at each country's actual 2019 real GDP per capita. The "Gap (%)" columns report the percentage gap of the fixed-adaptation path below the risk-dependent path; positive numbers indicate that ignoring risk-dependent adaptation understates 2050/2100 income. The "World aggregate" row reports population-weighted (2019 population fixed) means across the same set of cyclone-affected countries. Real GDP per capita is in thousands of constant 2015 USD.

	2050				2100			
Statistic	SSP5 base.	Fixed adapt.	Risk-dep. adapt.	Gap (%)	SSP5 base.	Fixed adapt.	Risk-dep. adapt.	Gap (%)
Mean (across countries)	34.4	31.4	31.6	0.5	106.8	85.1	86.9	1.8
S.D.	30.2	28.5	28.6	1.1	61.9	58.8	59.1	5.3
Median	22.7	20.7	20.8	0.0	89.0	64.5	67.1	0.9
P25	14.6	13.1	13.0	-0.0	65.9	48.1	49.4	-0.0
P75	43.3	40.8	41.7	0.6	122.3	97.6	100.7	3.2
World aggregate (pop-weighted)	31.3	25.3	25.5	0.6	87.7	54.7	55.7	1.8
Countries	57	57	57	—	57	57	57	—

Table 16 reports the cross-country distribution of projected 2050 and 2100 real GDP per capita under the three scenarios. Two patterns stand out. First, cyclones inflict economically meaningful gross damages over the projection horizon, even in the fixed-adaptation scenario: at the 2050 horizon, mean real GDP per capita is \$31,400 under fixed cyclone damages versus \$34,400 under the SSP5 baseline (a roughly 9% shortfall); by 2100 the gap widens to \$85,100 versus \$106,800, a roughly 20% shortfall. These numbers reflect the compounded effect of the baseline cyclone damage $\hat{\beta}_1\hat{\lambda}_{it}$ over eight decades.

Second, risk-dependent adaptation systematically reduces these damages. By 2050 the average country is 0.5% richer under the risk-dependent scenario than under the fixed-adaptation scenario, and by 2100 the average gap widens to 1.8%. Weighted by 2019 population, the world-aggregate gap is of similar order—0.6% by 2050 and 1.8% by 2100—reflecting that the populous low-to-middle-income economies (China, India, Bangladesh, Indonesia) sit roughly in the middle of the country-level adaptation-gain distribution. The cross-country distribution is highly skewed. Countries whose NHPP arrival rate has trended upward over 1980–2019 and is projected to continue rising—Nicaragua, Fiji, Singapore, Malaysia, Guatemala, Indonesia, Panama, Costa Rica—see 2100 adaptation gains in the range of 7% to 19% of real GDP per capita. The value of adaptation grows large over long horizons precisely because the future-risk deviation $\tilde{\lambda}_{it}$ widens further from the country's 1980 baseline as climate change progresses, and the adaptation offset $\hat{\beta}_2\tilde{\lambda}_{it}$ accumulates year after year.

For countries with falling or stable projected risk, the adaptation channel delivers little or even slightly negative gain. The model implies that countries whose risk has fallen below their 1980 baseline have correspondingly let adaptive capital depreciate, so when a cyclone does strike the per-disaster damage is somewhat larger than the constant-damage benchmark. Lao PDR, the Philippines, Japan, and Russia are visible examples in Figure 7. The aggregate cross-country mean therefore understates the value of adaptation for countries whose climate exposure is rising, which is exactly the set of countries where the integrated-assessment literature most needs forward-projectable estimates of adaptation gains.

The forward-projection exercise is, in our view, the central economic payoff of using a time-varying NHPP risk model instead of a time-invariant cumulative-exposure moderator. The static framework of Hsiang and Jina (2014) can document that high-exposure countries have adapted historically, but it cannot answer the policy question of how much additional adaptation a country will deploy as its risk rises above the historical range—because in that framework the moderator n_i is fixed at the country's historical realization. Our framework delivers exactly that forward calculation, country by country, by combining the estimated panel coefficient $\hat{\beta}_2$ with NHPP-projected risk paths. Under standard climate-economy growth assumptions, the

Figure 7: Projected real GDP per capita 2020–2100 and country-level adaptation gain in 2100

The left panel shows the cross-country mean projected real GDP per capita from 2019 to 2100 under three scenarios: the SSP5 baseline (black dotted; no climate damages), SSP5 plus expected cyclone damages with fixed adaptation $\hat{\beta}_1\hat{\lambda}_{it}$ (red), and SSP5 plus expected cyclone damages with risk-dependent adaptation $(\hat{\beta}_1 + \hat{\beta}_2\tilde{\lambda}_{it})\hat{\lambda}_{it}$ (blue). The gap between the blue and red lines is the cumulative adaptation gain at each future year. The right panel shows the country-level 2100 adaptation gain, $(GDP^C - GDP^B)/GDP^C$ in percent, sorted from largest to smallest. Countries with projected $\tilde{\lambda}_{it} > 0$ throughout 2020–2100 benefit from adaptation; countries with $\tilde{\lambda}_{it} < 0$ (NHPP-projected risk below their 1980 baseline, such as Lao PDR, Philippines, Japan) show a negative gain, reflecting the model's implication that adaptive capital has depreciated as risk has fallen.

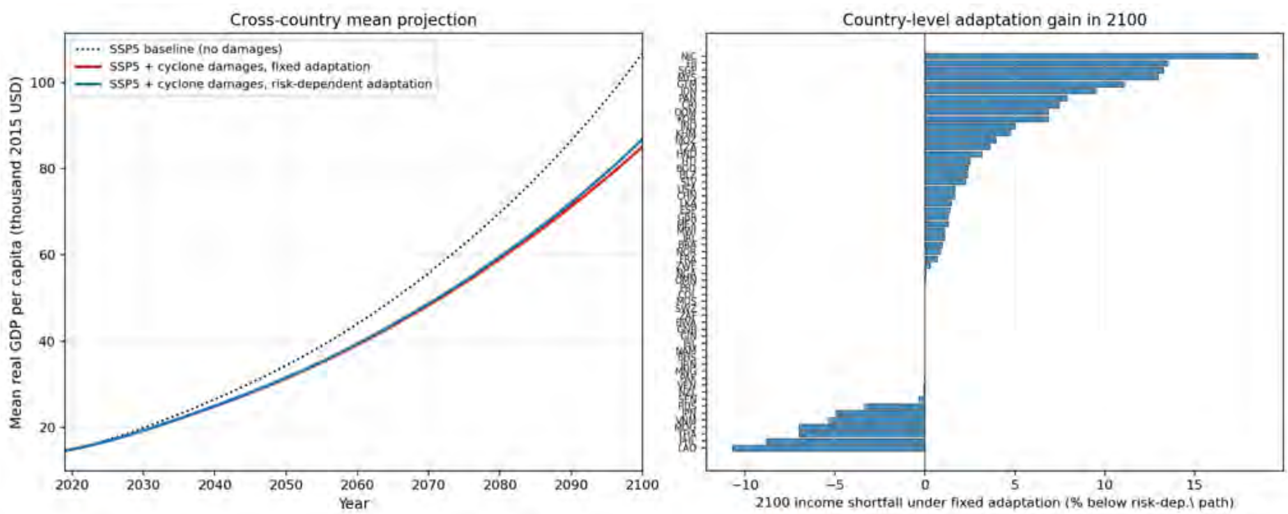

value of adaptation accruing to the most climate-stressed economies over the next eighty years is on the order of 10% to 20% of projected real GDP per capita.

7 Conclusion

This paper develops an empirical framework to measure the economic return to adaptation in a changing climate that exploits the time-inhomogeneity of disaster arrivals. We model each country's cyclone arrival process as a nonhomogeneous Poisson process with a quadratic time trend, then estimate in a panel regression with country and region-by-year fixed effects how the per-cyclone damage to GDP growth varies with the country's estimated time-varying risk. Cyclones inflict significantly smaller GDP-growth damages in country-years where risk is elevated relative to the country's 1980 baseline, consistent with optimizing agents spending more on adaption as the underlying disaster risk rises. The empirical interaction coefficient survives a battery of identification checks: lagged-disaster controls, country-specific trends, placebo permutations of disasters and risk, a horse race against the time-invariant cumulative-exposure moderator used in the prior literature, and a TWFE-robust binned dose-response estimator in the spirit of the recent difference-in-differences literature.

The framework's central payoff is forward-looking. Because our adaptation moderator is derived from a statistical model of the disaster arrival process, it can be projected forward under alternative climate scenarios—a calculation that is not available within frameworks whose adaptation moderator is fixed at historical realizations. Combining the estimated interaction coefficient with NHPP-projected risk paths and SSP5 growth assumptions, the value of risk-dependent adaptation over 2020–2100 reaches two to three percent of mean real GDP per capita on average across cyclone-affected countries and ten to twenty percent for the most climate-stressed economies. These magnitudes are inputs that integrated-assessment models, climate-related fiscal planning, and the pricing of catastrophe risk all require, and that the static cross-country adaptation literature has not been able to provide.

References

Auffhammer, M. (2022): "Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption," Journal of Environmental Economics and Management, 114, 102669.

Bakkensen, L. A., and R. O. Mendelsohn (2016): "Risk and Adaptation: Evidence from Global Hurricane Damages and Fatalities," Journal of the Association of Environmental and Resource Economists, 3(3), 555–587.

Barnett, M., W. Brock, and L. P. Hansen (2020): "Pricing uncertainty induced by climate change," The Review of Financial Studies, 33(3), 1024–1066.

Barreca, A., K. Clay, O. Deschenes, M. Greenstone, and J. S. Shapiro (2016): "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century," Journal of Political Economy, 124(1), 105–159.

Bilal, A., and D. R. Känzig (2024): "The macroeconomic impact of climate change: Global vs. local temperature," Discussion paper, National Bureau of Economic Research.

Borusyak, K., X. Jaravel, and J. Spiess (2024): "Revisiting Event-Study Designs: Robust and Efficient Estimation," Review of Economic Studies, 91(6), 3253–3285.

Bouwer, L. M., R. P. Crompton, E. Faust, P. Höppe, and R. A. Pielke Jr (2007): "Confronting disaster losses," Science, 318(5851), 753–753.

Burke, M., S. M. Hsiang, and E. Miguel (2015): "Global non-linear effect of temperature on economic production," Nature, 527(7577), 235–239.

Callaway, B., A. Goodman-Bacon, and P. H. C. Sant'Anna (2024): "Difference-in-Differences with a Continuous Treatment," NBER Working Paper No. 32117.

Carleton, T., A. Jina, M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, R. E. Kopp, K. E. McCusker, I. Nath, J. Rising, A. Rode, H. K. Seo, A. Viaene, J. Yuan, and A. T. Zhang (2022): "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits," The Quarterly Journal of Economics, 137(4), 2037–2105.

Cattaneo, M. D., R. K. Crump, M. H. Farrell, and Y. Feng (2024): "On Binscatter," American Economic Review, 114(5), 1488–1514.

Daley, D. J., and D. Vere-Jones (2003): "An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods," Springer.

de Chaisemartin, C., and X. D'Haultfœuille (2020): "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," American Economic Review, 110(9), 2964–2996.

Dell, M., B. F. Jones, and B. A. Olken (2012): "Temperature shocks and economic growth: Evidence from the last half century," American Economic Journal: Macroeconomics, 4(3), 66–95.

Dell, M., B. F. Jones, and B. A. Olken (2014): "What Do We Learn from the Weather? The New Climate-Economy Literature," Journal of Economic Literature, 52(3), 740–98.

Deryugina, T. (2017): "The Fiscal Cost of Hurricanes: Disaster Aid versus Social Insurance," American Economic Journal: Economic Policy, 9(3), 168–198.

Deschênes, O., and M. Greenstone (2007): "The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather," American Economic Review, 97(1), 354–385.

Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2014): "Optimal taxes on fossil fuel in general equilibrium," Econometrica, 82(1), 41–88.

Gourio, F., and C. Fries (2020): "Adaptation and the Cost of Rising Temperature for the U.S. Economy," Working paper, Federal Reserve Bank of Chicago.

Hong, H., N. Wang, and J. Yang (2023): "Mitigating Disaster Risks in the Age of Climate Change," Econometrica, 91(5), 1763–1802.

Hsiang, S. M., and A. S. Jina (2014): "The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence From 6,700 Cyclones," Working paper, National Bureau of Economic Research.

Jordà, Ò. (2005): "Estimation and Inference of Impulse Responses by Local Projections," American Economic Review, 95(1), 161–182.

Kriegler, E., N. Bauer, A. Popp, F. Humpenöder, M. Leimbach, J. Strefler, L. Baumstark, B. L. Bodirsky, J. Hilaire, D. Klein, et al. (2017): "Fossil-fueled development (SSP5): An energy and resource intensive scenario for the 21st century," Global Environmental Change, 42, 297–315.

Kutoyants, Y. A. (1998): "Statistical Inference for Spatial Poisson Processes," Springer-Verlag.

MacKinnon, J. G., and M. D. Webb (2018): "The wild bootstrap for few (treated) clusters," The Econometrics Journal, 21(2), 114–135.

National Academies of Sciences, Engineering, and Medicine (2016): Attribution of extreme weather events in the context of climate change. National Academies Press.

Newell, R. G., B. C. Prest, and S. E. Sexton (2021): "The GDP-temperature relationship: Implications for climate change damages," Journal of Environmental Economics and Management, 108, 102445.

Nordhaus, W. D. (2017): "Revisiting the social cost of carbon," Proceedings of the National Academy of Sciences, 114(7), 1518–1523.

Politis, D. N., and J. P. Romano (1994): "The Stationary Bootstrap," Journal of the American Statistical Association, 89(428), 1303–1313.

Schlenker, W., and M. J. Roberts (2009): "Nonlinear temperature effects indicate severe damages to US crop yields under climate change," Proceedings of the National Academy of sciences, 106(37), 15594–15598.

Sun, L., and S. Abraham (2021): "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," Journal of Econometrics, 225(2), 175–199.

Vuong, Q. H. (1989): "Likelihood ratio tests for model selection and non-nested hypotheses," Econometrica, 57(2), 307–333.