Natural Disasters, Adaptation and Default Risk

Carlos Bolivar
University of Minnesota

Federico Dueñas
International Monetary Fund

Oscar M. Valencia
Inter-American Development Bank

September 9, 2026*

Abstract

This paper studies how sovereign default risk shapes public investment in climate adaptation. We build a sovereign default model in which governments invest in adaptation to reduce disaster damages. Sovereign risk depresses adaptation through two mechanisms: higher borrowing costs and debt overhang under long-term debt. These effects generate feedback between risk and disaster exposure. Using a new panel of adaptation spending from budget records, we show that after catastrophic disasters spreads rise where prior adaptation was low, and that losses are smaller where adaptation budgets recently increased. In our calibration, adaptation yields the largest welfare gains near the default boundary, where financing is tightest.

Keywords: Sovereign debt, Default risk, Climate adaptation, Natural disasters
JEL classification: F34, F41, H63, Q54.

*We thank the editor and an anonymous referee for comments that substantially improved the paper.
E-mails: carlos.i.bolivar.h@gmail.com, f.duenas@uniandes.edu.co and oscarva@iadb.org.

1 Introduction

The effect of climate change on the frequency and severity of natural disasters has become a central concern for policymakers, particularly in emerging markets where fiscal constraints limit governments' ability to respond.[1] Beyond their immediate human cost, natural disasters destroy productive capacity, reduce governments' ability to repay debt, and increase sovereign default risk. As climate change intensifies these risks, lenders demand higher interest rates, further constraining the fiscal space of vulnerable countries. In this context, a natural policy response is *adaptation*: public investment that reduces the economic losses caused by natural disasters, such as improved drainage systems, early warning systems, or pre-positioned emergency supplies.

This paper studies the optimal adaptation policy of a government that borrows in international markets and faces the risk of both default and natural disasters. We develop a quantitative sovereign default model in the tradition of Arellano (2008), augmented with the risk of natural disasters as in Mallucci (2020). The main innovation is to model adaptation as an endogenous government investment that reduces economic losses during disasters. In the model, adaptation has three effects on fiscal sustainability: it increases government resources in disaster states, it improves the conditional distribution of future income when output is persistent, and it reduces the risk of default, which is reflected in higher sovereign bond prices.

Adaptation is different from standard public investment in capital. Building a road or a factory increases productivity in good and bad times. Adaptation, by contrast, pays off when natural disasters occur: it prevents collapse in bad times. In this sense, adaptation is an investment in insurance. Because it protects the economy precisely when default risk is highest, after a disaster, adaptation affects bond prices through a different channel than standard investment: a unit of adaptation spending reduces the probability of default because its benefits are concentrated where they matter most to creditors.

The first contribution of the paper is empirical. We build a panel of adaptation spending from government budgets, identifying the agencies whose mandate is protection for 19 emerging and developing economies over two decades. With this measure, we document two patterns. Extreme droughts have a lower impact on economic growth and spreads where the adaptation budget has recently increased. And after catastrophic disasters, sovereign spreads rise only where past adaptation spending was low.

Next, we present a model of adaptation policy. The central mechanism is that sovereign risk constrains adaptation, creating a feedback loop between sovereign risk and higher exposure to

[1] See Mallucci (2020) for an analysis of the effect of natural disasters on the risk of default.

natural disasters. As default risk rises, borrowing becomes more expensive, so the government scales back protection and the economy remains exposed. Moreover, in default episodes, with cutoff market access, the government must finance adaptation entirely from domestic resources, and as a result exposure to natural disasters increases. This feedback is strongest at intermediate debt levels, where adaptation can determine whether the government survives a natural disaster without defaulting on its debt. On average, access to adaptation is worth 1.20 percent of consumption and is substantially more in these states.

The second key theoretical result is that the adaptation investment is inefficiently low. With long-term debt and limited commitment, adaptation suffers from debt overhang (Aguiar et al., 2009; Hatchondo et al., 2016): some of the marginal benefit of higher resilience is capitalized in the price of outstanding bonds, so part of the return accrues to legacy bondholders rather than to the government at the time it invests. We make this distortion explicit by deriving the corresponding wedge in the government's Generalized Euler Equation and quantifying it.

This inefficiency motivates climate adaptation-linked debt. We introduce a green bond whose coupon steps up when investment falls below a pre-specified rule, modeled on the step-up architecture of the sustainability-linked issues of Chile and Uruguay in 2022. We let the government choose its mix of green and standard debt. At step-ups in the range observed in markets, the government voluntarily holds green bonds for about two thirds of its debt, follows the rule in most states, and gains 0.063 percent of consumption. The government's use of the instrument reflects a trade-off. Issuing more green bonds increases resources today because investors pay a premium for the higher payments they expect under step-ups. This choice also moves the price of both bonds through two mechanisms. Step-ups raise future payments in repayment states, which can increase default risk and lower bond prices. However, the covenant increases adaptation, reduces the exposure to disaster, and reduces the risk of future default, which can increase the prices of bonds.

Related Literature. Our paper belongs to the quantitative literature on sovereign default based on the seminal contributions of Arellano (2008) and Aguiar and Gopinath (2006); see Aguiar and Amador (2021) for a recent survey. Within this literature, our work is most closely related to papers studying the interaction between default risk and investment. Aguiar et al. (2009) show that the debt burden of sovereign debt depresses investment: a government that cannot commit to not tax capital ex post discourages private investment when its debt is high, so investment falls with the debt burden. With long-term bonds, Hatchondo et al. (2016) show that a debt policy that raises the price of outstanding debt gives part of its return to legacy bondholders, which the government does not internalize. Esquivel (2024) formalizes

a related result for private capital and documents misallocation under sovereign risk, and Esquivel and Samano (2023) studies fiscal consolidation with capital accumulation. We build on this strand by showing that the same debt overhang logic applies to adaptation investment, with a specific policy implication: governments remain under-prepared for natural disasters because the fiscal architecture of long-term debt does not reward resilience.

Our paper also belongs to the literature on natural disasters and climate risk. This literature distinguishes between ex-ante adaptation, prevention investments such as climate-resilient infrastructure, insurance strategies, and contingent financing, and ex-post adaptation, which encompasses emergency responses and post-disaster reconstruction, including international aid and multilateral financing. Dell et al. (2012) document that temperature shocks reduce growth in poor countries, and the evidence on disasters suggests that countries with lower ex-ante adaptation capacity experience greater economic losses, worsening fiscal sustainability, and increasing sovereign risk. Hsiang and Narita (2012) find that adaptive capacity can be expensive with declining marginal benefits, which is particularly relevant for tropical cyclones where adaptation efforts show limited returns. Despite these findings, the sovereign debt literature has largely treated adaptation as an exogenous factor rather than a strategic government decision, and has not studied how the government's fiscal position constrains its ability to invest.

A growing body of work explores the link between climate risk and sovereign borrowing costs. Cevik and Jalles (2022) find that countries with greater climate resilience face lower sovereign spreads. Boitan and Marchewka-Bartkowiak (2022) document that weaker climate management capacity is associated with higher borrowing costs in Europe, particularly in the southern and southeastern regions. Beirne et al. (2021) show that climate vulnerability has a substantially larger effect on bond yields than resilience has on reducing them. Gómez González et al. (2024) highlight the asymmetric nature of climate risk in sovereign financing costs, showing that countries with higher spreads experience more pronounced effects. Kling et al. (2025) estimate that climate-vulnerable developing countries paid more than 100 basis points in additional borrowing costs between 2007 and 2016. On the default side, Klomp (2017) finds that large-scale natural disasters increase the probability of sovereign default by approximately three percentage points. These papers document the correlation between climate vulnerability and borrowing costs, but do not model the feedback: that higher borrowing costs themselves discourage adaptation, which increases vulnerability further.

Most directly, our paper builds on general equilibrium models linking sovereign risk and natural disasters. Mallucci (2020) extends a sovereign default model with long-term bonds to incorporate disaster risk, showing that increased climate shocks increase borrowing costs and restrict market access, but does not model adaptation as an endogenous policy choice. Phan

and Schwartzman (2023) incorporate capital accumulation and financial frictions, finding that climate shocks have persistent growth effects with recovery delays of more than two decades, but similarly treat adaptation as exogenous. Barucci et al. (2025) model sovereign default with climate risk and let the country invest in pollution abatement, but focus on debt renegotiation rather than the interaction between debt structure and adaptation incentives. Grosse-Steffen et al. (2025) study fiscal rules in a sovereign default model where natural disasters are the key shock, showing that rules with escape clauses create fiscal space for post-disaster spending, but do not endogenize adaptation investment. Also related is Arellano et al. (2023a), who study debt relief policies in a sovereign default model with an epidemiological model of COVID-19; we consider adaptation policy in an analogous setting with climate shocks. We address the gap in this literature by allowing the government to choose adaptation investment, which enables us to study the feedback loop between sovereign risk and disaster vulnerability and to show that the countries most exposed to climate shocks are precisely the ones least able to invest in preparedness.

Outline. The remainder of the paper is organized as follows. Section 2 measures the adaptation in the data. Section 3 describes the model and Section 4 characterizes the optimal policy and the commitment wedges. Section 5.1 discusses the calibration. Section 5 presents the quantitative results, including the green-bond economy; Section 7 reports the sensitivity to maturity and the strength of the agreement, and Section 8 concludes.

2 Adaptation in the Data

This section measures adaptation spending and tests whether it protects countries against natural disasters. Section 2.1 describes the data. Section 2.2 studies droughts. We start with droughts because they are the one hazard that we observe for every country and year before it causes any damage and independently of the country's protection (Vicente-Serrano et al., 2010). We ask whether a drought costs less output where past adaptation spending is larger and whether sovereign spreads rise less after a drought begins where that spending is larger. Section 2.3 extends the analysis of spreads to every type of catastrophic event, those with recorded damage above the 75th percentile of the sample.

2.1 Measurement

Adaptation We measure the investment in adaptation from government budgets. The source is the World Bank BOOST database, which harmonizes line-level public expenditure

records drawn from national treasury systems and reports executed spending by administrative unit. Within each budget, we identify the agencies whose mandate is adaptation, the environment ministries and their technical services, the institutes of water and hydraulic infrastructure, and the civil-protection and disaster-risk offices. We measure adaptation as the share of total executed spending by these agencies. The resulting panel covers 19 emerging and developing economies in Latin America, Africa, Eastern Europe, and Asia between 2003 and 2024. Table 1 reports descriptive statistics. The median adaptation share is 0.59 percent of government spending, with an interquartile range from 0.19 to 1.14 percent; environment agencies account for roughly three quarters of identified amounts and water infrastructure for most of the remainder.

Table 1: The Adaptation Budget Panel: Descriptive Statistics

Adaptation share of government spending (percent)	
Mean	0.79
Median	0.59
Interquartile range	0.19–1.14
Minimum, maximum	0.05, 4.49
Median within-country coefficient of variation	0.32
Composition of identified adaptation spending (percent)	
Environment agencies and technical services	73.6
Water and hydraulic infrastructure	25.3
Civil protection and disaster-risk offices	1.1

Note: Executed spending of adaptation-mandated agencies as a share of total executed government spending, identified from the administrative classification in the World Bank BOOST database and matched by agency name within each country's own budget structure. Countries: Albania, Brazil, Burkina Faso, Burundi, Chile, Croatia, Dominican Republic, Guatemala, Haiti, Kenya, Mauritania, Mexico, Moldova, Poland, Senegal, Seychelles, Solomon Islands, Ukraine, and Uruguay.

Disaster risk Disaster data comes from two sources. EM-DAT reports on direct economic deaths and losses from natural disasters. Second, the SPEI database reports a monthly drought index on a global grid of half a degree for 1901–2024 (Vicente-Serrano et al., 2010). The index is the net rainfall of evaporation over the previous 12 months.The drought index of a country-year is minus one times the lowest monthly value in the year, so that a higher value means that the country's driest twelve months of that year were drier.

Output Real GDP growth is the log change of chain-linked real GDP from the Global Macro Database, in percent, for the budget-panel countries over 2004–2023.

Spreads Daily sovereign spreads come from the J.P. Morgan EMBI Global for 25 emerging markets. We work with the spread log and report changes in log points, so that a change of ten is ten percent of the spread.

2.2 Droughts

Output losses We relate real GDP growth to two variables: the drought index and the country's past adaptation spending. Given a country's climate, whether a year turns out drier than usual is random, so once country and year effects are removed, we argue that the index does not depend on the country's protection.[2] The adaptation measure is the mean of the budget share over the five years before t, standardized within country, so that one unit is one standard deviation of the country's own history. This is the panel design of Dell et al. (2012), with past adaptation spending as a trait that moderates the effect of the weather. We estimate

$$g_{it} = \beta_1 \,\text{Drought}_{it} + \beta_2 \,\text{Drought}_{it} \times A_{i,t-1} + \beta_3 \, A_{i,t-1} + \alpha_i + \lambda_t + \epsilon_{it}, \tag{1}$$

where g_{it} is the percent growth and $A_{i,t-1}$ is the adaptation measure. The prediction is $\beta_2 > 0$: a drought costs less growth where past spending is higher. Table 2 reports the estimates for the 19 budget-panel countries over 2004–2020, 227 country-years.

Table 2: Droughts, Past Adaptation Spending, and Output Growth

	(1) Baseline	(2) Drought × income	(3) Two-year growth
Drought	−0.60	−0.68	−0.56
	(0.147)	(0.104)	(0.291)
Drought × adaptation	0.58	0.60	1.39
	(0.064)	(0.046)	(0.002)
Adaptation	0.21	0.13	0.45
	(0.233)	(0.513)	(0.182)
Observations	227	227	218

Notes: Real GDP growth in percent on the drought index, its interaction with past adaptation spending (mean budget share over the five previous years, standardized within country), and that spending, with country and year fixed effects; 19 countries, 2004–2023. Column (2) adds the index interacted with lagged log GDP per capita, column (3) uses two-year growth, column (4) adds country trends. Two-sided p-values from standard errors clustered by country in parentheses.

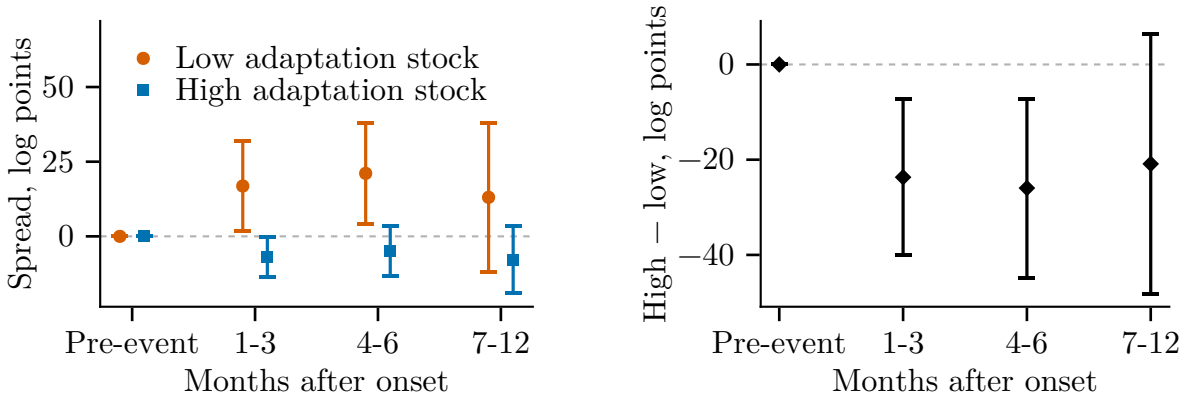
[2]Appendix B gives the construction of the index and the alternative measures we pre-specified.

At the country's usual level of past spending, a drought costs growth: $\beta_1 = -0.60$ percentage points per unit of the index, so a year at the 90th percentile of the index instead of the median lowers growth by 0.56 percentage points. The interaction is positive $\beta_2 = 0.58$ ($p = 0.064$; wild cluster bootstrap with one side $p = 0.052$): where past spending is one standard deviation above the country's usual level, the drought cost disappears. The estimate is similar when the drought is allowed to interact with income (column 2) and larger over two years of growth (column 3).

Spreads Next, we analyze the effect on spreads. The monthly index dates droughts, which allows an event study on sovereign spreads with events selected on the weather. An event is a drought onset defined as 0 is a normal year for that place, −1 is drier than 84% of the years recorded there, and −1.5, the threshold for a severe drought, is drier than 93%. We averaged the index over each country's grid cells, weighting by area, which gives one value per country and month.

For each event, we take the monthly log spread relative to its mean over the six months before the onset and subtract the same quantity averaged over the EMBI countries with no onset within twelve months of that calendar month. Events whose window overlaps the 2008–2009 financial crisis, the March 2020 repricing, a war on or at the country's border, or a sovereign restructuring are excluded; Appendix B states the rules and lists the events.

Figure 1: Sovereign Spreads around Drought Onsets



(a) Group means

(b) High minus low

Note: Ten drought onsets in six budget-panel countries with daily EMBI coverage, 2005–2023, that reach a severe drought within twelve months, split by whether the country's past adaptation spending before the onset is below or above the sample median. Each point is the cross-event mean of the event-level average spread over the months shown, relative to the event's own months −6 to −1 and net of the contemporaneous change in control countries; whiskers are 90 percent confidence intervals from cross-event variation.

This yields ten onsets in six budget-panel countries between 2005 and 2023: two each in the Dominican Republic, Mexico, Poland, and Ukraine, and one each in Brazil and Uruguay. We split them by the country's adaptation measure before the onset, at the sample median. Figure 1 shows the paths. Where past spending is low, spreads increase by 17 log points in months 1–3 after the onset and by 21 in months 4–6; where it is high, they do not rise. The difference is −24 log points in months 1–3 and −26 in months 4–6, with exact randomization p-values of 0.028 and 0.032 over the 252 possible splits of ten events into two groups of five, and is negative when any one event is dropped.

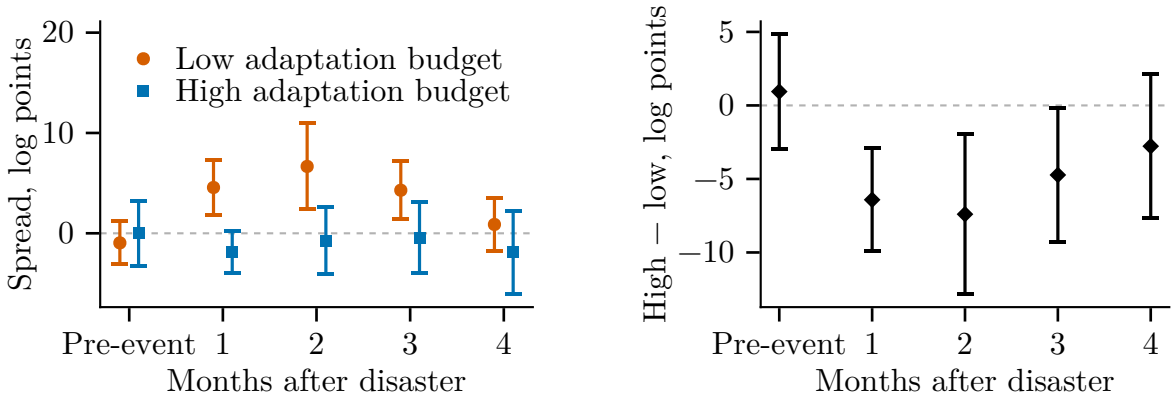
2.3 Catastrophic Disasters and Spreads

We now extend the spread analysis to every type of disaster. We restrict attention to catastrophic disasters, following the sample definition of Cavallo et al. (2013): they find no effect of typical disasters on output and at most a fragile one for the very largest, so if spreads respond to disasters at all, it is these events that should show it. Events are the country-years of budget-panel countries with daily EMBI coverage whose EM-DAT climate damages exceed the 75th percentile of the sample. For each event, we compute the daily log spread relative to its own mean over trading days −10 to −3, net of the same change in the spreads of emerging markets with no catastrophic disaster in the window, average it within months of 21 trading days, and then average across events. Events whose window overlaps the 2008–2009 financial crisis or the March 2020 restriction are excluded, as in Section 2.2.

This yields eight events in five countries between 2010 and 2023: three in Mexico, two in Guatemala, and one each in Brazil, Chile, and Poland. We split them by the country's lagged adaptation share relative to the sample median. Because events are selected on realized damages, which adaptation itself reduces, a damage-selected event in a country with a large adaptation budget must have been physically harsher to clear the threshold; this selection works against finding a smaller spread response where adaptation is high, not for it.

Figure 2 summarizes the results. In the pre-event window, high- and low-adaptation countries look similar: the mean difference is 0.9 log points (s.e. 2.4), and the rank correlation between the lagged adaptation share and pre-event drift is statistically insignificant (with the sign running against the post-event pattern). After the disaster, spreads move only in the low-adaptation group. Relative to its own pre-event mean, the low-share group rises by 4.6 log points in month 1, 6.7 in month 2, and 4.3 in month 3, before drifting back toward zero; the high-share group stays near its pre-event level throughout.

Figure 2: Sovereign Spreads around Catastrophic Disasters



(a) Group means

(b) High minus low

Note: Eight climate disasters in budget-panel countries with daily EMBI coverage, 2010–2023, defined as events with EM-DAT damages above the sample 75th percentile and dated from event-level EM-DAT records. Events are split by whether the country's lagged adaptation budget share is below or above the sample median. Each point is the cross-event mean of the event-level average spread in that month, relative to the event's own pre-window mean and net of the contemporaneous change in control countries; whiskers are 90 percent confidence intervals from cross-event variation.

In addition, we test the pattern using Fisher-style randomization inference (Fisher, 1935; Conley and Taber, 2011; Athey and Imbens, 2017; Young, 2019): we reassign the spending ranks between events and count how often the pattern is as strong as the one observed. The four high-spending events saw spreads rise 6.2 log points less than the four low-spending events; only 4.3 percent of the 70 possible splits of eight events into two groups of four produce a gap as large as our partition.

The second test uses the order of the events, not just the split. We rank the eight events by spending and, separately, by the rise in spreads, and check whether the two rankings run in opposite directions: whether the biggest spenders sit at the bottom of the response list. They do: the rank correlation is -0.55, and only 8.6 percent of the 40,320 possible orderings of the spending ranks across events produce a correlation this negative ($p = 0.086$, one-sided). These are permutation tail areas with events as units, and the eight events come from five countries, so responses within a country may be correlated. Treating countries as units instead, the same rank test on country means gives a correlation of -0.60 with $p = 0.18$ over the 120 orderings, and dropping any one country leaves the correlation negative, with p between 0.07 and 0.20. The pattern is therefore a regularity consistent with protection, robust in sign, and modest in significance.This result is consistent with the asymmetric spread response documented by Gómez González et al. (2024) through protection: the same shock

increases spreads only where the past adaptation is low.

Robustness test The budget measure is our preferred proxy for adaptation, but it does not cover the countries where the most destructive storms landfall. Appendix B therefore repeats the design in an independent event sample: tropical cyclones with landfall winds of at least 125 knots (category 4 on the Saffir–Simpson scale). Among countries with daily sovereign-spread data, this yields 15 events between 2006 and 2025. Landfall dates and wind speeds come from IBTrACS. Finally, as BOOST is not available for these countries, we split countries using the IMF environmental-expenditure measure because the budget panel does not extend to the cyclone belt.

The pattern is consistent. Two months after landfall, spreads are 3.8 log points above their pre-event level in low-spending countries, but are back at their pre-event level in high-spending countries. The implied difference is -4.5 logarithmic points, with a one-sided p-value of 0.050. The two groups experience storms of comparable physical severity.

In the next section, we propose a model that helps us to understand the interaction of investment in adaptation and sovereign risk and evaluate the effect of green bonds in this context.

3 Model

We study a small open economy in which the government borrows in international markets. As in Eaton and Gersovitz (1981), the government lacks commitment, and each period decides whether to default on its debt. Following Mallucci (2020), we assume that the economy experiences natural disasters that affect output with an exogenous probability. The main innovation is that the government can invest in adaptation, a technology that reduces the economic impact of natural disasters.

Timing. Within each period, the sequence of events is as follows. The government enters with debt b, an adaptation stock a, and income y, which already incorporates the effect of any current disaster given the predetermined stock. Then it observes the default cost ν and decides whether to default. After the default decision, it chooses the adaptation investment i, which determines the next period stock a', and, if in good standing, the new debt b'.

3.1 Endowment

In each period, the government receives a stochastic endowment of tradable goods y_t. The government's income follows an AR(1) process subject to two shocks: a standard endowment shock ϵ_t that captures fluctuations in the business cycle and a disaster shock that captures the possibility of a natural disaster.

Let $\xi_t \in \{0, 1\}$ be an indicator variable drawn from a Bernoulli distribution with parameter π, where $\xi_t = 1$ denotes the occurrence of a disaster. Conditional on a disaster occurring, the intensity γ_t is drawn from an i.i.d. distribution $G(\gamma)$, independent of all other shocks.

When a disaster strikes, the economy suffers an output cost that depends on the intensity of the disaster and the economic adaptability stock $a_t \geq 0$:

$$h(a_t, \gamma_t) = \gamma_t \exp\left(-a_t \, \gamma_t / \bar{\gamma}\right) \tag{2}$$

Without adaptation, the economy bears the full cost γ_t, and each unit of the stock reduces the damages proportionally at the rate $\gamma_t/\bar{\gamma}$, where $\bar{\gamma}$ is the mean intensity; this fixes the units in which the stock is measured. Protection is therefore proportionally strongest for the most severe disasters.For a given stock, income follows an AR(1) in logs hit by the proportional disaster loss:

$$\log \hat{y}_t = \rho \log y_{t-1} + \mu + \epsilon_t, \qquad y_t = \hat{y}_t \left[1 - \xi_t \, h(a_t, \gamma_t)\right], \tag{3}$$

where $\hat{y}_t$ is income before the current disaster, $\epsilon_t \sim N(0, \sigma_\epsilon^2)$, and μ sets the mean of $\hat{y}_t$ to one. Because income is persistent, a disaster has dynamic effects: a negative shock to y_t propagates into future periods. Adaptation therefore has both a static benefit, reducing the immediate output loss, and a dynamic benefit, improving the conditional distribution of future income. This persistence amplifies the value of adaptation in a setting where disaster shocks are purely transitory.

3.2 Government

The government has time-separable preferences with a CRRA utility function , and discounts the future at a rate β:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [u(c_t) - d_t \nu_t] \tag{4}$$

Here $d_t \in \{0, 1\}$ is an indicator that equals one if the government defaults. The term ν_t

is a stochastic cost of default drawn from a distribution $f(\nu)$ independent of all other state variables. This stochastic default cost serves two purposes: it makes the default decision, and with it the expected continuation value, a smooth function of the state, which allows us to derive clean analytical characterizations of the default threshold; and it provides a tractable computational device for smoothing the default decision in the numerical solution, the extreme-value device used by Dvorkin et al. (2021).

After the default decision, the government chooses adaptation investment $i_t \geq 0$, which accumulates into a stock that depreciates at rate δ_a:

$$a_{t+1} = (1 - \delta_a)a_t + i_t, \qquad i_t \geq 0 \tag{5}$$

Investment is irreversible: the government cannot consume the installed stock. The cost of investment is

$$f(i_t) = p_a \, i_t \tag{6}$$

where p_a is the unit price of adaptation. The government borrows using long-term bonds, following Hatchondo and Martinez (2009), Chatterjee and Eyigungor (2012), and Arellano and Ramanarayanan (2012). A bond issued in period t promises to pay $\kappa(1-\delta)^{j-1}$ units of the tradable good in period $t+j$ for all $j \geq 1$, where δ is the rate at which the bond amortizes and κ is the coupon.Denoting by b_t the outstanding face value of the debt and n_t the net issuance of new bonds:

$$b_{t+1} = b_t(1 - \delta) + n_t \tag{7}$$

The budget constraint of the government depends on its standing in international markets. Under repayment, the government services its debt and can issue new bonds at price q_t:

$$c_t + f(i_t) = y_t - \kappa b_t + q_t(b_{t+1} - (1 - \delta)b_t) \tag{8}$$

Under default, all debt is forgiven but the government loses access to international markets and suffers an output cost. Following Chatterjee and Eyigungor (2012), the output cost is a nonlinear function of income, so the effective endowment under default is:

$$y_t^D = y_t - \max\{0, \eta_1 y_t + \eta_2 y_t^2\} \tag{9}$$

The budget constraint under default is:

$$c_t + f(i_t) = y_t^D \tag{10}$$

Although excluded from international markets, the government regains access to each period with probability θ.

3.3 International Investors

As in Eaton and Gersovitz (1981), there is a continuum of identical, risk-neutral lenders with access to a risk-free asset at the gross interest rate R. Lenders have sufficient resources to buy any amount of government bonds. The break-even condition for bond prices is the following:

$$q_t = \mathbb{E}\left[\frac{(1-d_{t+1})[\kappa + (1-\delta)q_{t+1}]}{R}\right] \tag{11}$$

3.4 Competitive Equilibrium

Definition 1. (Competitive Equilibrium) Given an initial condition $\{b_0, a_0, y_0\}$, an equilibrium consists of a sequence of prices $\{q_t\}$ and government policies $\{b_t, a_t, d_t\}$ such that

(i) Given prices, $\{b_t, a_t, d_t\}$ solves the government problem (4) subject to (8) and (10);

(ii) Given government policies, q_t satisfies the break-even condition (11).

4 Optimal Policy

This section presents our main analytical results. We establish three findings. First, a lower adaptation increases default incentives, formalizing the link between adaptation and sovereign risk. Second, under long-term bonds, the government's lack of commitment opens two wedges relative to the efficient allocation, one on borrowing and one on adaptation investment.Third, the government's fiscal position shapes its ability to invest in adaptation, creating a feedback loop between sovereign risk and disaster vulnerability.

4.1 Government Problem and Default Decision

The government lacks commitment, so we focus on the Markov equilibrium. The relevant state for the government is the level of debt b, the adaptation stock a, the income y and the cost of default ν. We summarize the state with $\chi = \{b, a, y, \nu\}$. Through its investment

choice, the government also shapes the distribution of future income y' through equations (5) and (3). The government's problem is the following.

$$V(\chi) = \max_{d \in \{0,1\}} (1-d)\, V^R(b,a,y) + d\,[V^D(a,y) - \nu] \tag{12}$$

Here, $V^R(b,a,y)$ is the value of repayment and $V^D(a,y) - \nu$ the value of default. The cost ν is paid in the period of default and affects no choice, so we write the default value as $V^D(a,y)$, which does not depend on ν, less the realized cost. The value of repayment is

$$\begin{aligned} V^R(b,a,y) &= \max_{i \geq 0,\, b'} u(c^R) + \beta \mathbb{E}[V(\chi')] \qquad (13) \\ &\text{subject to} \\ c^R &= y - \kappa b - f(i) + q(y,b',a')(b' - (1-\delta)b) \\ a' &= (1-\delta_a)a + i \\ y'(y,a',\gamma') &= \rho y - \xi' h(a',\gamma') + \epsilon' \end{aligned}$$

The government induces a distribution on the endowment over the next period by choosing the stock a'. In particular, it chooses the economy's exposure to natural disasters, which is equivalent to choosing a conditional distribution of future income. In addition, adaptation has a precautionary motive as in Aguiar et al. (2009) and Esquivel (2024) as the government can accumulate adaptation in good times and reduce it during crises or natural disasters. The default value is

$$\begin{aligned} V^D(a,y) &= \max_{i \geq 0} u(c^D) + \beta \mathbb{E}[\theta V(0,a',y',\nu') + (1-\theta) V^D(a',y')] \qquad (14) \\ &\text{subject to} \\ c^D &= y^D - f(i) \\ a' &= (1-\delta_a)a + i \\ y'(y,a',\gamma') &= \rho y - \xi' h(a',\gamma') + \epsilon' \end{aligned}$$

Although excluded, the government invests out of its own resources and incurs no further cost. We characterize the default decision by defining the threshold default cost that makes the government indifferent between repayment and default.

$$\hat{\nu}(b, a, y) = V^D(a, y) - V^R(b, a, y) \tag{15}$$

The government defaults if and only if the realized default cost is below this threshold:

$$d_t(\chi) = \begin{cases} 1 & \text{if} \quad \hat{\nu}(b, a, y) > \nu_t, \\ 0 & \text{otherwise.} \end{cases} \tag{16}$$

The conditional probability of default for a given endowment, debt, and adaptation stock is

$$F(\hat{\nu}(b, a, y)) = \int_{\underline{\nu}}^{\hat{\nu}(b,a,y)} f(\nu) d\nu \tag{17}$$

Adaptation and the default threshold. A larger stock in place when a disaster strikes leaves more income, and a government is tempted to default when income is low. The cleanest statement takes the current disaster to be the last shock the economy faces.[3]

Lemma 1. *Let income be constant after the current period, $y_{t+s} = \bar{y}$ for all $s \geq 1$, exclusion after default permanent, $\theta = 0$, and the default decision deterministic, $\nu \equiv 0$. The default sets then shrink with the stock in place when the disaster strikes: if the government defaults on a stock a_H, it defaults on every stock $a_L < a_H$.*

Proof. See Appendix E.1. □

The proof is closely related Arellano (2008). With no further disasters, neither regime invests, so government carry the same stock forward under default and repayment, and the repaying government's option to default tomorrow makes its continuation at least the defaulter's. A government that defaults therefore cannot be receiving net resources from the market, so its consumption under repayment is below its consumption under default. Also, concavity then makes an extra unit of income worth more under repayment, so the threshold falls with income; and a larger stock in place leaves more income.

In general, the result rests on one property: a government with more resources is less tempted to default. Whenever it holds, a larger stock lowers the threshold when a natural disaster hits, and investment in adaptation shrinks the default region. With disaster risk and persistent income, the comparison of their continuations is no longer available; Appendix E.1 provides sufficient conditions. The lemma 1 establishes the first direction of the relationship

[3]The default cost has a continuous density, so the expected continuation value is smooth in the state; we take the differentiability of V^R and V^D in the state as given, as is standard.

between adaptation and sovereign risk: countries with low adaptation face a higher risk of default and, consequently, lower rates of bonds. Section 4.3 studies the other direction: how sovereign risk constrains adaptation investment.

Price function. Using the break-even condition of international investors and the probability of default, the equilibrium price function is:

$$q(y, b', a') = \mathbb{E}\left[\frac{(1 - d(\chi'))[\kappa + (1-\delta)q(y', b'', a'')]}{R}\right] \tag{18}$$

Definition 2. (Markov Equilibrium) A Markov perfect equilibrium is defined by a set of policy functions $\{\mathcal{B}, \mathcal{A}, \mathcal{D}\}$, a value function V, and a price function $q(y, b', a')$ such that:

(i) Given the price function, $\{\mathcal{B}, \mathcal{A}, \mathcal{D}\}$ solves the government's problem in every state and V achieves the maximum.

(ii) $q(y, b', a')$ satisfies the break-even condition of foreign investors (18).

4.2 Commitment Wedges in Adaptation and Borrowing

When the government is in good standing and the investment is interior, the necessary conditions for optimality are[4]:

$$u'(c^R)[q_{b'}(b' - (1-\delta)b) + q] = \beta\mathbb{E}[(1 - d(\chi'))u'(c^{R'})(\kappa + (1-\delta)q')] \tag{19}$$

$$u'(c^R)[f'(i) - q_{a'}(b' - (1-\delta)b)] = \beta\pi\mathbb{E}\left[-h_a(a', \gamma')\,\mathcal{V}_y(\chi') \,|\, \xi' = 1\right] + \beta(1-\delta_a)\mathbb{E}\left[u'(c')f'(i')\right] \tag{20}$$

where $\mathcal{V}_y(\chi') = (1 - d(\chi'))\,\partial_y V^R(b', a', y') + d(\chi')\,\partial_y V^D(a', y')$ is the marginal value of income in the next period. By the envelope theorem it satisfies

$$\partial_y V^R(b, a, y) = u'(c^R)\left[1 + q_y\left(b' - (1-\delta)b\right)\right] + \beta\rho\,\mathbb{E}\big[\mathcal{V}_y(\chi')\big], \tag{21}$$

[4]The stochastic default cost smooths the expected continuation value $\mathbb{E}_{\nu'}[V(\chi')]$ in (b', a'): the default set is $\{\nu' < \hat{\nu}(b', a', y')\}$ and its probability moves continuously with the choices, taking the differentiability of V^R and V^D in the state as given.

$$\partial_y V^D(a,y) = u'(c^D)\left[1 - D'(y)\right] + \beta\rho\,\mathbb{E}\left[\theta\,\partial_y V(0,a',y',\nu') + (1-\theta)\,\partial_y V^D(a',y')\right], \quad (22)$$

where q_y is the derivative of the price schedule with respect to income and $D(y)$ the output cost of default. Equation (19) is the standard Euler equation for debt, equating the marginal benefit of increasing current consumption with the marginal cost of repaying in the next period. Equation (20) is the Euler equation for adaptation investment. On the left-hand side, $f'(i)$ is the marginal cost of investment, while $q_{a'}(b' - (1-\delta)b)$ captures the offset effect of higher bond prices due to higher adaptation; equation (24) below signs it. On the right-hand side, we have the effect of adaptation in the next period. The first term is the protection benefit: $-h_a(a',\gamma')$ is the marginal damage reduction, weighted by the probability of a disaster and by the marginal value of the income it saves. That value exceeds the marginal utility of consumption for two reasons that (21)–(22) make explicit: income is persistent, so a unit of output saved today raises expected income in every future period, and under repayment a unit of income is also sold into a better price. The second term comes from the higher stock of adaptation: a fraction $(1-\delta_a)$ of the stock survives into the next period, where it saves the future government its own marginal investment cost.

The economic logic of the potential inefficiency is as follows. When the government increases adaptation, it reduces the probability of default in case of natural disaster. This raises the value of all outstanding long-term bonds, not just newly issued ones. Part of this value increase is a capital gain for holders of previously issued debt, legacy bondholders, that the government does not internalize. The same margin operates on borrowing itself: issuing additional debt tomorrow dilutes today's bondholders as in Hatchondo and Martinez (2009).

The Generalized Euler Equation (GEE). To formalize this result, we rewrite the price function as:

$$q(y,b',a') = \mathbb{E}\left[\frac{(1-d(y',b',a'))[\kappa + (1-\delta)q(y',\mathcal{B}(b',a'),\mathcal{A}(b',a'))]}{R}\right] \quad (23)$$

where $b'' = \mathcal{B}(b',a')$ and $a'' = \mathcal{A}(b',a')$ are the equilibrium policy functions iterated one period forward. Taking the derivative of the price function with respect to adaptation and using the relevant envelope conditions yields:

$$q_{a'} = \mathbb{E}\left[\frac{1}{R}\Big(-f(\hat{\nu}')\left[\xi'\big(-h_a(a',\gamma')\big)\,\partial_y\hat{\nu}' + \partial_a\hat{\nu}'\right][\kappa + (1-\delta)q''\right]$$

$$+(1-F(\hat{\nu}'))(1-\delta)\left[\xi'\left(-h_a(a',\gamma')\right)q''_y+q''_{b'}\mathcal{B}_a+q''_{a'}\mathcal{A}_a\right]\Big)\Bigg] \quad (24)$$

where $\hat{\nu}'=\hat{\nu}(b',a',y')$ is next period's default threshold, $q''=q(y',\mathcal{B},\mathcal{A})$ is next period's price with q''_y its derivative with respect to income, and $\mathcal{B}_a$ and $\mathcal{A}_a$ are the total derivatives of next period's policies with respect to a', including their response through y'. The first term in (24) is the change in the default probability: a larger stock moves the threshold through post-disaster income, $\xi'(-h_a)\,\partial_y\hat{\nu}'$, and through the state directly, $\partial_a\hat{\nu}'$, the two channels of Lemma 1. When a government with more resources is less tempted to default, the bracket is negative and the term positive: a larger stock leaves the government more resources when a disaster strikes, and the default probability falls. The second term collects the effects on next period's price: through income, which raises the price when the government is less likely to default, and through the government's future borrowing and adaptation choices, $\mathcal{B}_a$ and $\mathcal{A}_a$, whose signs the theory does not fix. The second term exists only because debt is long-term: it is scaled by $(1-\delta)$ and vanishes with one-period bonds. Figure 3 verifies that the total is positive in the states where default risk is present.

The derivative of the price function depends on its own future values, creating a recursive structure. Following Arellano et al. (2023b), we use equations (19) and (20) to solve for $q_{b'}$ and $q_{a'}$ and define[5]:

$$H_{b'}(b',a')=\frac{\beta\,\mathbb{E}\left[(1-d(\chi'))u'(c^{R'})(\kappa+(1-\delta)q')\right]}{u'(c^R)\,(b'-(1-\delta)b)}-\frac{q}{b'-(1-\delta)b} \quad (25)$$

$$H_{a'}(b',a')=\frac{f'(i)}{b'-(1-\delta)b}-\frac{\beta}{u'(c^R)\,(b'-(1-\delta)b)} \quad (26)$$
$$\times\left(\pi\,\mathbb{E}\left[\left(-h_a(a',\gamma')\right)\mathcal{V}_y(\chi')\,\middle|\,\xi'=1\right]+(1-\delta_a)\,\mathbb{E}\left[u'(c')f'(i')\right]\right)$$

Substituting $H_{b'}$ and $H_{a'}$, evaluated at next period's state, for $q''_{b'}$ and $q''_{a'}$ in (24) and combining with (20) yields the Generalized Euler Equation for adaptation:

$$u'(c^R)\Bigg(f'(i)-(b'-(1-\delta)b)\,\mathbb{E}\Bigg[\frac{1}{R}\Big(-f(\hat{\nu}')\left[\xi'\left(-h_a(a',\gamma')\right)\partial_y\hat{\nu}'+\partial_a\hat{\nu}'\right][\kappa+(1-\delta)q'']$$
$$+(1-F(\hat{\nu}'))(1-\delta)\left[\xi'\left(-h_a(a',\gamma')\right)q''_y+H_{b'}\mathcal{B}_a+H_{a'}\mathcal{A}_a\right]\Big)\Bigg]\Bigg)$$

[5]See Mateos-Planas and Ríos-Rull (2015) for a general discussion of the GEE in sovereign debt models.

$$= \beta\pi\, \mathbb{E}\big[\big(-h_a(a',\gamma')\big)\, \mathcal{V}_y(\chi') \,\big|\, \xi'=1\big] + \beta(1-\delta_a)\, \mathbb{E}\big[u'(c')f'(i')\big] \quad (27)$$

The GEE accounts for the cost of the government's lack of commitment. Because the government cannot commit to future policies, it takes as given the future policy functions $\mathcal{A}$ and $\mathcal{B}$. Under commitment, the government would choose future adaptation and borrowing so that $H_{b'}$ and $H_{a'}$ in equations (25) and (26) equal zero in every period: the envelope conditions would hold and future policies would not distort current prices. Without commitment, future governments deviate from this plan, and the GEE contains a wedge that reflects how costly these deviations are. Two properties of the wedge matter for what follows. It is the sum of a borrowing component, $H_{b'}\mathcal{B}_a$, and an adaptation component, $H_{a'}\mathcal{A}_a$, which need not be of comparable size; and it is scaled by $(1-\delta)$, so it vanishes as debt becomes short term, since with one-period bonds there is no legacy stock to revalue. Section 7 measures both components.

4.3 Sovereign Risk and the Cost of Adaptation

The previous section established that adaptation reduces sovereign risk. We now study the other direction: how sovereign risk shapes the government's ability to invest in adaptation. This question is largely absent from the climate risk literature, which typically studies adaptation as a decision constrained by costs and technology but not by the government's fiscal position. In our model, the cost of adaptation depends fundamentally on how the government finances it, and the cost of financing depends on sovereign risk.

Adaptation financed with debt. When the government has access to international markets, it can finance adaptation by issuing new debt rather than reducing consumption. To isolate this channel, consider the set of policies (b', a') that are consistent with a fixed level of current consumption $\bar{c}$:

$$\bar{c} = y - \kappa b - f(i) + q(y, b', a')(b' - (1-\delta)b) \quad (28)$$

Along this locus, the government trades off higher future debt against greater adaptation. The net marginal benefit of this trade-off depends on how bond prices respond to both debt and adaptation. When the government increases adaptation, the default probability falls (Lemma 1) and, unless the policy terms in (24) reverse it, bond prices rise, partially offsetting the cost of additional borrowing; the calibrated model confirms that $q_{a'} > 0$ in the states where default risk is present (Figure 3). This mechanism is related to Esquivel and Samano (2023), who show that a fiscal consolidation which lowers sovereign risk raises

private investment and output, so that a policy that improves the price of debt pays for part of itself; here the price response to adaptation plays that role.

In Appendix E, we derive the full optimality condition for debt-financed adaptation and decompose its benefits into two components: the direct income effect of reducing disaster losses, which depends on the covariance between marginal utility and disaster intensity; and the insurance effect, which depends on how adaptation changes the payoff differential between repayment and default states.

Adaptation financed from domestic resources. When the government has defaulted and is excluded from international markets, debt financing is unavailable. The first-order condition for adaptation under default is:

$$\begin{aligned} u'(c^D) f'(i) = \beta\pi\theta\, \mathbb{E}\big[\big(-h_a(a',\gamma')\big)\, \partial_y V(0,a',y',\nu') \,\big|\, \xi'=1\big] \\ + \beta\pi(1-\theta)\, \mathbb{E}\big[\big(-h_a(a',\gamma')\big)\, \partial_y V^D(a',y') \,\big|\, \xi'=1\big] + \beta(1-\delta_a)\, \mathbb{E}\big[u'(c')f'(i')\big] \end{aligned} \tag{29}$$

The left-hand side shows that every unit of adaptation must be paid for by reducing consumption—there is no possibility of spreading the cost over time through borrowing. The right-hand side values the income the stock protects with the marginal values of income of (21)–(22), weighted by whether the government has regained access. It reveals that the benefit of adaptation depends on how long the government expects to remain excluded. If the probability of regaining market access (θ) is low, the government expects to remain in autarky with low consumption and high marginal utility, which increases the incentive to invest in adaptation. As θ rises, the government anticipates being able to borrow again soon, which raises expected future consumption, lowers marginal utility, and reduces the incentive to invest.

The feedback loop. Comparing the repayment and default cases reveals the feedback between sovereign risk and adaptation. When the government has access to credit on favorable terms, it can finance adaptation with cheap debt, reducing the risk of disasters without sacrificing consumption. As sovereign risk increases, borrowing becomes more expensive, the government cuts adaptation investment, and the economy becomes more vulnerable to disasters, which further increases sovereign risk. In the extreme case of default, the government loses entirely the option of debt-financed adaptation, making investment costly. This feedback implies that the countries most exposed to natural disasters, which tend to be the ones with higher sovereign risk, are precisely the countries least able to invest in

adaptation.

5 Quantitative Analysis

In this section, we present the quantitative analysis. We first describe the calibration and report the calibrated parameters and the model's fit. We then analyze the model's implications for bond prices, optimal policy, and welfare, and close with the two experiments that answer the paper's normative questions: the role of debt maturity and the value of commitment through adaptation-linked contracts.

5.1 Calibration

The calibration maps the empirical regularities of Section 2 to the model's parameters.

Disaster Process Disasters arrive with probability $\pi = 0.109$, the annual frequency of catastrophic climate events in the EM-DAT panel of emerging markets. Conditional on a disaster, the intensity γ is discretized to four values spanning 5.0 to 12.9 percent of output, with 98 percent of the conditional mass on the smallest; the scale is set so that the conditional mean loss is 5.1 percent of output, the EM-DAT mean for catastrophic events. The bin probabilities are the relative frequencies of predicted damages in the readiness-based damage regression of Appendix C; only the frequency π and the conditional mean are taken from the moments used in Section 2. The distribution is therefore close to a point mass with a thin upper tail, and the results below should be read as describing a typical catastrophic event rather than the full range of disaster severities.

Adaptation Technology Two parameters describe the adaptation technology, and each is identified by one moment of the budget data. The unit price p_a is set to the coefficient of variation of the adaptation share (32.3 percent in BOOST), which measures how much the budget swings around its own mean. The depreciation rate δ_a is set to the lag-1 autocorrelation of that share (0.66): a stock that depreciates slowly generates a persistent spending series, and a fast-depreciating one must be rebought every year. Together these two moments say that adaptation budgets are volatile and only moderately persistent, which is what pins the technology down.

We fix the adjustment cost at $\iota = 0$ rather than calibrating it. With the depreciation rate that the persistence moment selects, a positive ι does not buy additional persistence: it lowers the level of investment and, at high δ_a, collapses the stock, so the two parameters cannot be

separately identified from the spending series. Setting it to zero makes the investment cost linear, which is also the case in which the optimality conditions of Section 4 are exact.

The spending level is not a target. The measured 0.167 percent of GDP covers recurrent items only, so it is a lower bound on adaptation spending rather than an estimate of it; we require the model to clear it and it does, at 1.00 percent. The damage coefficient of the regression of recorded losses on the budget share, equation (31) in Appendix A, is likewise not targeted. It is too imprecisely estimated to discipline a parameter, and we report it in Table 4 as a validation moment instead.

Debt Block The bond pays a geometric coupon with decay $\delta = 0.0824$, the value used by Mallucci (2020); it implies an expected life of $1/\delta = 12.1$ years and a Macaulay duration of 8.2 years at the risk-free rate, and the coupon is $\kappa = \delta + r$ so that a risk-free bond prices at par. The spread of the model is $\kappa/q_t - \delta - r$. The discount factor $\beta = 0.91280$ and the default cost curvature $\eta_2 = 0.41740$ jointly match a mean spread of 2.0 percent and a mean external-debt-to-GDP ratio of 45 percent, round values based on the median EMBI spread and government debt ratio of the budget-panel economies with traded debt over 2000–2023 (2.3 and 47 percent; debt from the fiscal space database of Kose et al., 2022). The remaining parameters are standard and reported in Table 3.

5.2 Calibrated Parameters and Fit

Table 3 reports the parameter values. The coefficient of relative risk aversion ($\sigma = 2$), the international interest rate ($R = 1.0451$), and the probability of regaining market access ($\theta = 0.33$) are standard in the literature. The persistence of income ($\rho = 0.92$), the standard deviation of the income innovation ($\sigma_\epsilon = 0.046$), and the linear default cost parameter ($\eta_1 = -0.198$) follow Mallucci (2020); income is discretized in ten Tauchen nodes that span the standard deviations 2.5. The stochastic default cost has the logistic distribution implied by the extreme-value smoothing of Appendix D, with scale 10^{-4}. The four remaining parameters are calibrated as described in Sub-section 5.1: (β, η_2) jointly to the mean spread and the debt ratio, and (p_a, δ_a) to the coefficient of variation and the lag-1 autocorrelation of the adaptation share. The spending level is imposed as a lower bound rather than a target, and the damage coefficient is not a target at all.

Table 3: Parameters

	Value	Source
Standard in the literature		
Risk aversion	$\sigma = 2$	Arellano (2008)
Interest rate	$R = 1.0451$	Arellano (2008)
Probability of regaining access	$\theta = 0.33$	Arellano (2008)
Income persistence	$\rho = 0.92$	Mallucci (2020)
Income innovation, s.d.	$\sigma_\epsilon = 0.046$	Mallucci (2020)
Output loss, linear	$\eta_1 = -0.198$	Mallucci (2020)
Maturity decay	$\delta = 0.0824$	Mallucci (2020)
Disaster process, from EM-DAT		
Probability of disaster	$\pi = 0.109$	Catastrophic event frequency
Mean disaster cost	$\bar{\gamma} = 0.051$	Mean damages, share of GDP
Calibrated internally		
Discount factor	$\beta = 0.91280$	Mean spread
Output loss, quadratic	$\eta_2 = 0.41740$	Mean debt-to-GDP
Unit price of adaptation	$p_a = 0.01424$	Spending volatility (c.v.)
Depreciation of adaptation	$\delta_a = 0.75$	Spending autocorrelation

Table 4 reports the fit. The four targeted moments are closely matched: the spread within two and a half basis points, the debt ratio within a tenth of a percentage point of GDP, the volatility of spending within a fifth of a percentage point and the persistence of the share of adaptation within 0.018. The model's spending level, 1.00 percent of GDP, sits above the measured 0.167 percent, as it should, since the measured level covers recurrent items only.

Table 4: Model Fit

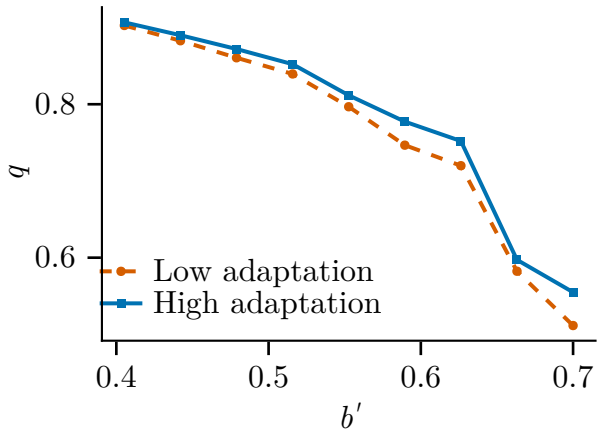
	Data	Model
Targeted moments		
Mean spread (%)	2.0	2.02
Mean debt-to-GDP (%)	45	45.1
Spending coefficient of variation (%)	32.3	32.1
Adaptation share, autocorrelation	0.66	0.68

Notes: Spread and debt targets are round values near the medians for budget-panel economies with traded debt (EMBI spreads; fiscal space database of Kose et al., 2022). The adaptation spending level is Antigua and Barbuda's recurrent adaptation budget, 0.167 percent of GDP (hand-collected); the coefficient of variation is from BOOST. We set four parameters to match four moments one-for-one.

5.3 Bond Prices and Adaptation

Figure 3 plots the bond price schedule as a function of new borrowing in a state with income just below its mean and no disaster active, for a low and a high level of the adaptation stock on which the issue is priced. The schedules decrease in borrowing, and adaptation shifts them up. The price response is concentrated in the borrowing region where the risk of default is high, the state dependence discussed in Section 4.2.

Figure 3: Price Function and Adaptation



Note: Bond price schedule against new borrowing b', at income one grid point below the ergodic mean with no disaster active. The entering adaptation stock is held at the ergodic mean, so the two schedules differ only in the stock a' the issue is priced on.

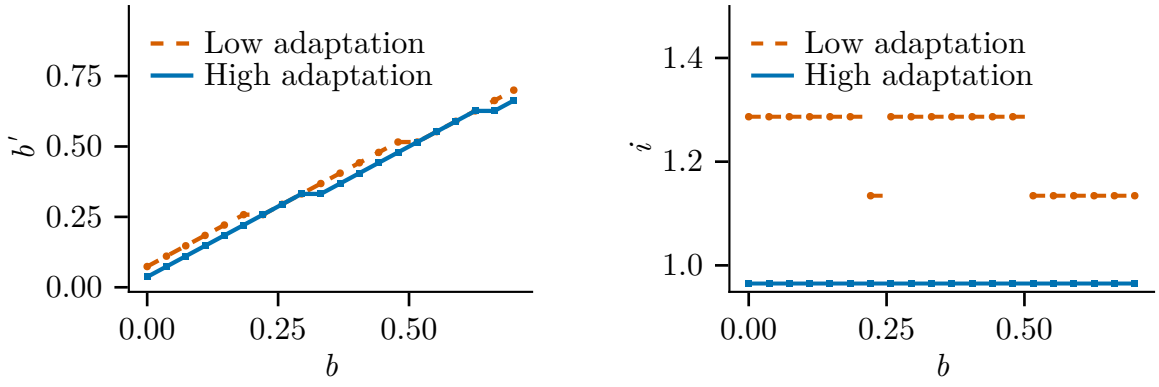
5.4 Optimal Policy

Figure 4 shows the optimal debt and adaptation policies when the economy is faced with a severe natural disaster and income is above its mean. We plot policies for two initial adaptation stocks: no stock, and the highest stock the economy reaches in the ergodic distribution. At this state the government repays at every debt level on the grid, so both panels are on-equilibrium throughout.

The panel 4a shows the debt policy. The exposed government borrows weakly more than the protected one at every level of initial debt, and strictly more over most of the range. The wedge is one point of the debt grid wherever it opens, and it opens at thirteen of the twenty debt levels plotted. The government with no protection carries the full disaster loss into its budget, so absorbing it without borrowing would require a large drop in consumption; the protected government, whose stock already absorbs much of that loss, does not need the

extra resources and rolls its debt over instead.

Figure 4: Optimal Policy



(a) Debt

(b) Adaptation investment

Note: Policy functions against initial debt b, at income three grid points above the mean with the most severe disaster active, for initial adaptation stocks $a = 0$ and $a = 1.29$, the largest stock reached in the ergodic distribution. Panel (b) plots investment $i = a' - (1 - \delta_a)a$.

The panel 4b shows the investment policy. The government invests at every debt level, but how much depends on the initial stock. The protected government invests exactly its depreciation, holding its stock at the top of the ergodic range: it has already built the protection it wants and simply maintains it. The exposed government invests more than that at every debt level, building the stock it lacks within a single period: over most of the debt range it reaches the protected government's stock outright, and at the highest debt levels it stops one grid point short. It sustains that investment even there, where the resources have to come from a new issue rather than from income, and lenders accommodate it because the price of the issue depends on the adaptation it funds. This is the debt-financed adaptation channel of Section 4.3.

5.5 Welfare

We measure the welfare benefits of adaptation by computing, the percentage increase in consumption across all future histories that a household would require to be indifferent between an economy where the government can invest in adaptation and one where it cannot. The values we compare are the policy-evaluation values of Appendix D: the expected discounted utility of the consumption stream generated by the equilibrium policies at the

equilibrium prices. Utility is homogeneous in degree $1-\sigma$ in consumption. The equivalent consumption gain λ in the state χ therefore satisfies

$$(1+\lambda(\chi))^{(1-\sigma)}\, V^{NA}(\chi) = V(\chi)$$

where $V^{NA}(\chi)$ is the value when the adaptation investment is not available. Figure 5 shows the welfare gains as a function of initial debt, averaged across income levels with each level weighted by its share of simulated periods, separately for states with and without an active disaster.

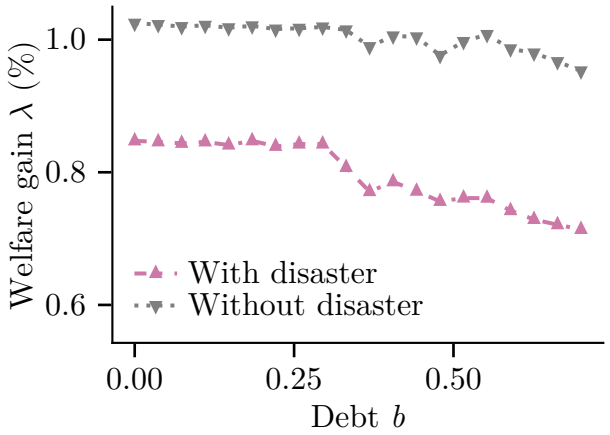


Figure 5: Welfare Gains from Adaptation Investment

Note: Consumption-equivalent welfare gain from access to adaptation investment, computed at zero initial adaptation ($a = 0$), against initial debt. Each line averages the ten income levels, weighted by their shares of simulated periods; the two lines condition on the most severe disaster being active and on no disaster.

The average consumption-equivalent gain is 1.20 percent, evaluated along the ergodic path; the figure instead holds the initial stock at zero, which is why the levels plotted are lower. The welfare gain from adaptation falls with debt because adaptation and default insure the disaster losses. When debt is low, the government pays back in disasters, so consumption drops to absorb the loss, and adaptation delivers the largest benefit for households. When debt is high, the government is closer to default, so part of the disaster loss is shifted to creditors, and there is less household consumption left to protect. Quantitatively, with a disaster active, the average gain decreases from 0.85 to 0.71 percent of consumption as debt increases.

At high income, the gain rises again as debt approaches the default boundary, where additional investment raises bond prices enough to keep market access: it peaks at 1.26

percent of consumption at a debt of 0.63. At low income, no such rise appears and the gain is largest at zero debt, at 0.95 percent. Both figures are for the income-split version of the calculation, which underlies the average plotted in Figure 5.

6 Green Bonds

Governments are increasingly using green bonds. We focus on bonds whose coupon increases when the issuer misses a stated target. In particular, we model a bond with a target on adaptation investment (as in the sovereign issues of Chile and Uruguay in 2022). In this section, we argue that with a large enough step-up green bonds could implement the investment allocation the government would choose under commitment.

6.1 Portfolio Choice

Structure of the bond A green bond pays the coupon κ plus a step-up ϕ in every period in which the adaptation investment falls below a rule $\bar{a}'(y, b, a)$. We set the rule to be the investment policy the government would choose if it could commit to the investment margin; Appendix F defines this commitment economy and how the rule is computed. We assume that the government trades both bonds in the international markets. Green and standard bonds are subject to default, and the government chooses the composition of its debt each period.

Price Write ϑ for the green share of the debt outstanding and ϑ' for the green share of the debt the government carries into next period, b'; the government trades both bonds every period, so its net issue of green bonds is $\vartheta' b' - (1-\delta)\vartheta b$, which can be negative. The state of this economy is the baseline state augmented by the composition, $\chi = \{b, a, y, \nu, \vartheta\}$. Bond prices are determined by cash flows, and the two bonds differ in exactly one cash flow, the step-up. Pricing that difference characterizes the green premium:

Lemma 2. *Suppose the two bonds pay the same base coupon κ, mature at the same rate δ, and are treated alike in default, and that the green bond pays an extra coupon $\phi > 0$ in every period in which investment falls below the rule. Then*

$$q_g(\chi) \;=\; q_s(\chi) \;+\; \mathbb{E}\left[\sum_{t=1}^{\infty} \frac{(1-\delta)^{t-1}}{R^t} \prod_{j=1}^{t} \big(1 - d(\chi_j)\big)\, \phi\, v(\chi_t) \;\middle|\; \chi_0 = \chi\right],$$

where $d(\chi)$ and $v(\chi)$ indicate default and violation of the rule. In particular $q_g(\chi) \geq q_s(\chi)$,

with equality if and only if $v(\chi_t) = 0$ *almost surely at every date* t *at which the government has not already defaulted.*

Proof. See Appendix E.1. □

The green bond trades above the standard bond by the expected present value of its incremental increases, collected only while the government repays.

The composition choice The government chooses the green share of the stock it carries forward by trading the premium against the burden. Consider the choice of ϑ' at an interior optimum, with b' and a' fixed. Write $n_g = \vartheta' b' - (1-\delta)\vartheta b$ and $n_s = (1-\vartheta')b' - (1-\delta)(1-\vartheta)b$ for the net issues of the two bonds, $n = n_g + n_s$ for the total net issuance, $P = q_g - q_s$ for the premium of Lemma 2, and $v(\chi) \in \{0, 1\}$ for a violation of the rule. The first-order condition is

$$u'(c)\, b'\, P \;+\; u'(c)\big[\, n\, \partial_{\vartheta'} q_s + n_g\, \partial_{\vartheta'} P \,\big] \;=\; \beta\, b'\, \mathbb{E}\Big[\big(1 - d(\chi')\big)\, u'(c^{R'})\, \big(\phi\, v(\chi') + (1-\delta)\, P(\chi')\big)\Big]. \tag{30}$$

Appendix E.3 derives it. The left-hand side is what a greener stock brings in today. The first term is the premium: by recomposing its stock, the government sells the step-up claim on the whole of b' and collects P per unit. The second term is the effect of the composition on the prices of today's net issues, to which we return below. The right-hand side is what a greener stock costs tomorrow. The next period's government inherits $\vartheta = \vartheta'$, and one more unit of green legacy costs it b' times the step-up it owes if it violates the rule, $\phi\, v(\chi')$, plus the premium at which it must refinance the surviving green bonds, $(1-\delta)P(\chi')$, which by Lemma 2 is the value of the step-ups they will owe later. Default wipes both, so the cost is paid only in repayment states, at the marginal utility of a government that is repaying.

Three features of the trade-off follow. First, the premium and the burden are the same stream of step-ups priced twice: lenders price it at the risk-free rate in P, the government at its own discount factor and marginal utility. Substituting the recursion of Lemma 2 for P, the premium and the right-hand side combine into $b'\,\mathbb{E}\big[(1-d(\chi'))\big(\phi v(\chi') + (1-\delta)P(\chi')\big)\big(u'(c)/R - \beta u'(c^{R'})\big)\big]$: the government gains from a greener stock when it values resources today more than lenders do, relative to the states in which the step-ups fall due. Those are states of high debt and low income, so the green bond is a way of borrowing against bad states, and it is most attractive when the government is constrained today. This is the pattern of Figure 8: the government holds green bonds at low income and light debt and lets them go as either rises.

Second, the composition changes default risk, and with it the prices of both bonds, through two channels that work against each other. A green legacy is a heavier liability for a repaying government and no liability at all for a defaulting one, so it raises the next period's default probability; this channel has an unambiguous sign because a government with a green legacy has weakly fewer resources at every choice than one without, while default treats the two alike. Against it, the covenant raises the investment in the next period, which lowers the default risk of the period afterward and raises the continuation price. The theory does not rank the two, and the sign of $\partial_{\vartheta'} q_s$ is a quantitative question. In the calibrated economy, the burden dominates: at the baseline maturity, the standard price is weakly decreasing in the green share in every state, for every step-up we solve (Figure 7), and the discipline channel prevails only in a minority of states when debt is short. The premium itself also falls with the green share $\partial_{\vartheta'} P \leq 0$, because a greener legacy induces more compliance and fewer escalations.

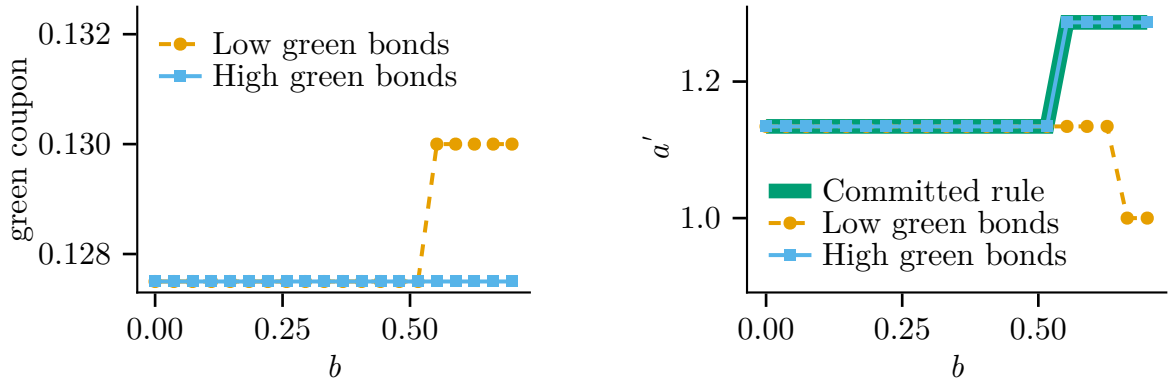
Third, one term is absent from (30). The composition revalues the legacy that lenders already hold, by $(1-\delta)\, b\, [\vartheta\, \partial_{\vartheta'} q_g + (1-\vartheta)\, \partial_{\vartheta'} q_s]$, and the government internalizes none of it: it is the composition analogue of the dilution wedge of Section 4.2, and its sing dependence of the effect of green issuance on current prices.

6.2 Quantification of Green Bonds

This subsection quantifies the effects of green bonds on adaption policy and welfare.

Coupons and adaptation policy We begin with Figure 6 showing the green bond at work and illustrating how it provides incentives to the government to implement the optimal investment in adaption.

When debt is low, the government can meet the rebuilding rule at no extra cost: it chooses the investment prescribed by the rules (panel b) and the green bond pays only the baseline coupon κ (panel a). As debt rises, the rule calls for higher adaptation than the government is willing to finance. It then under-invests relative to the rule (panel b), which triggers the covenant and raises the coupon by ϕ (panel a). In the stationary equilibrium, investment is below the rule in 11 percent of the periods in which the government repays.



(a) The green coupon

(b) The rule and the chosen investment

Figure 6: The Covenant at Work

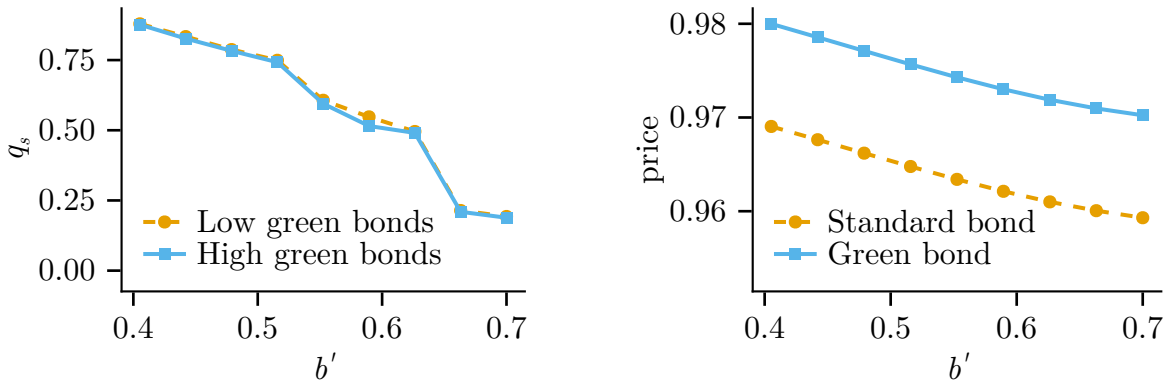
Note: Both panels at mean income with the mildest disaster active and an entering adaptation stock of 1.29, against debt, for a legacy stock that is all standard ($\vartheta = 0$) and all green ($\vartheta = 1$). Panel (a): the green bond's coupon, κ plus the step-up $\phi = 25$ basis points in periods in which the covenant is violated. Panel (b): the committed rule $\bar{a}'$, drawn as the thick band, with the chosen investment at each composition drawn on top of it.

The frequency with which the economy visits these states depends on the outstanding composition of the debt. The figure illustrates this mechanism: the government departs from the rule at a lower debt level when little of its debt is green, and keeps investing at the rule level to the top of the debt grid when most of it is. In counterfactual simulations that hold the debt composition fixed, a government constrained to an all-green stock complies 98 percent of the time, whereas one constrained to an all-standard stock violates the rule in 46 percent of periods, because with no green bonds outstanding, the step-up is not paid and the covenant has no bite. As a result, the covenant disciplines in direct proportion to the amount of green debt outstanding.

Prices Next, we describe how the composition of debt affects bond prices. Figure 7, illustrates the effect of green bonds on the price of standard bonds and also compares both prices.

Panel (a) plots the price of the standard bond for a debt stock carried forward that is all standard and for one that is all green. The two schedules nearly coincide. Where they separate, at intermediate and high debt, the all-green schedule lies weakly below the all-standard one, by at most a few cents on the dollar at this state. The composition moves the standard price only through default risk, and a stock of green debt changes that risk in the two ways of equation (30): the covenant raises adaptation, which lowers it, but a green

legacy is a costlier liability, through the step-ups it owes in violation states and the premium at which the surviving green bonds must be refinanced, both of which fall due exactly where the government is closest to default and both of which default wipes out, and that raises it. In the solved economy the second effect is the larger wherever the two differ, so a greener stock never raises the standard price; the composition matters for the standard price only at the margin. Panel (b) compares the green and standard prices in a state where government income is high. The green bond trades about 1% of face value above the standard bond, matching the expected present value of the step-ups a non-complying government is forecast to pay. In most states, however, the two prices are nearly the same: the government is below the rule in only 11 percent of repayment periods, and by Lemma 2 the premium is the expected value of future step-ups, so wherever a violation remains possible on some future path the premium is positive but small, and it vanishes only where the covenant is certain to be respected.



(a) Composition and the standard price

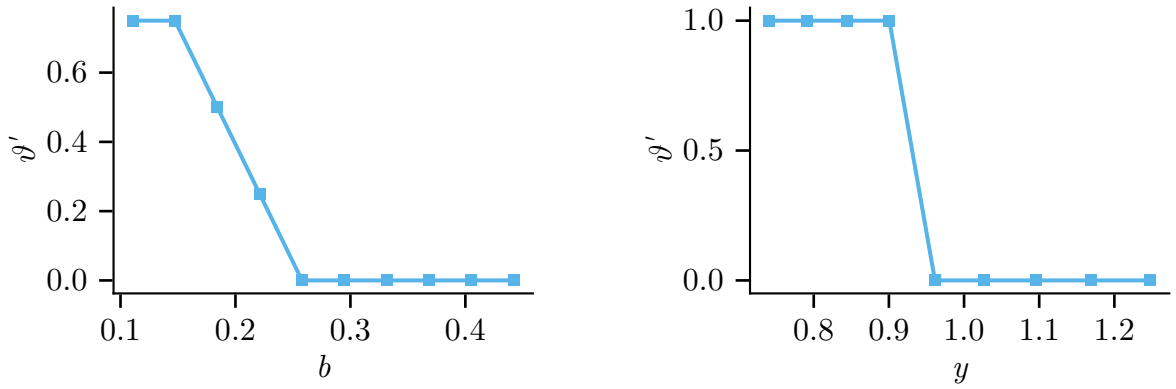
(b) Green and standard bond prices

Figure 7: The Price of Green Debt

Note: Panel (a): income two grid points below the ergodic mean, no disaster active, adaptation stock 0.53. Panel (b): top income, adaptation fixed at 0. Horizontal axis truncated at $b' = 0.4$ in both panels.

Portfolio choice Figure 8 illustrates how the government uses the new instrument. Panel (a) plots the chosen green share against debt. At low debt the government carries three quarters of its debt forward as green bonds; the share then steps down and reaches zero at a debt of 0.26, well short of its ergodic mean of 0.43, and stays there. Panel (b) plots the same choice against income: the share is the whole stock at the four lowest income levels and falls to zero above them. The two panels have different explanations. Against debt, the share turns down where default risk rises: near default the step-ups are unlikely to be collected,

so the premium a greener stock would earn is small, while the burden it would place on a government already close to the boundary is large. Against income, default risk falls as income rises, so it cannot be the reason; instead, at high income the government meets the rule on its own, the covenant never binds, and by Lemma 2 the green bond then carries no premium and buys no discipline the government lacks. The government therefore holds green bonds where the premium is worth most to it and the burden least: at low income, where it is constrained and the rule would otherwise be violated, and at light debt. The instrument's value runs the other way, and Figure 9 makes the divergence explicit.



(a) By debt

(b) By income

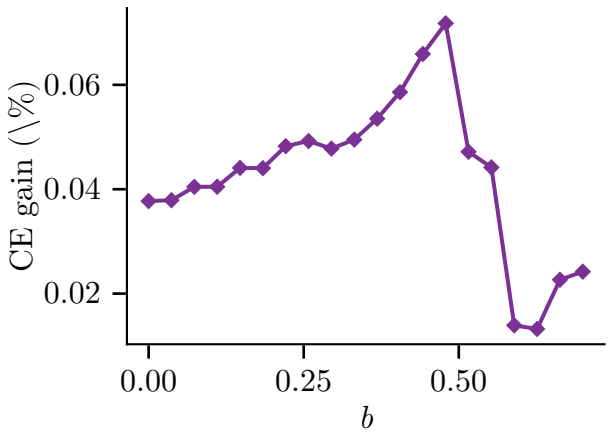
Figure 8: Green Bonds: the Chosen Composition

Note: Two-bond economy with a coupon step-up of $\phi = 25$ basis points per period of violation. Panel (a): the chosen green share of the debt carried forward, against debt, at income two grid points above the mean with no disaster active and an entering adaptation stock of 0.69, over $b \in [0.11, 0.44]$. Panel (b): the same choice against income, at middle debt with the most severe disaster active, omitting the top income grid point.

Welfare Figure 9 plots the value of having the instrument: the gain of the economy with the green bond over the economy without it, at each initial debt level. The gain is positive at every debt level. It rises with debt across the range the economy actually visits, from 0.038 percent of consumption at zero debt to a peak of 0.072 percent at a debt of 0.48, just under the ergodic 99th percentile, and falls beyond it to 0.024 percent at the top of the grid. The instrument is worth most where the commitment problem is worst, which is where debt is high and the temptation to cut adaptation is strongest. It is not, however, held most there: the composition policy in Figure 8 turns down once default risk is high, because near default the step-ups are unlikely to be collected, so the premium a greener stock would earn is small, while the burden it would place on a government already close to the boundary is large. The gap between where the bond is valuable and where it is held is the trade-off of equation (30)

in quantities. Near default, the step-ups stop being collected, the covenant loses its bite, and the gain fades.

Figure 9: Green Bonds: the Value of the Instrument

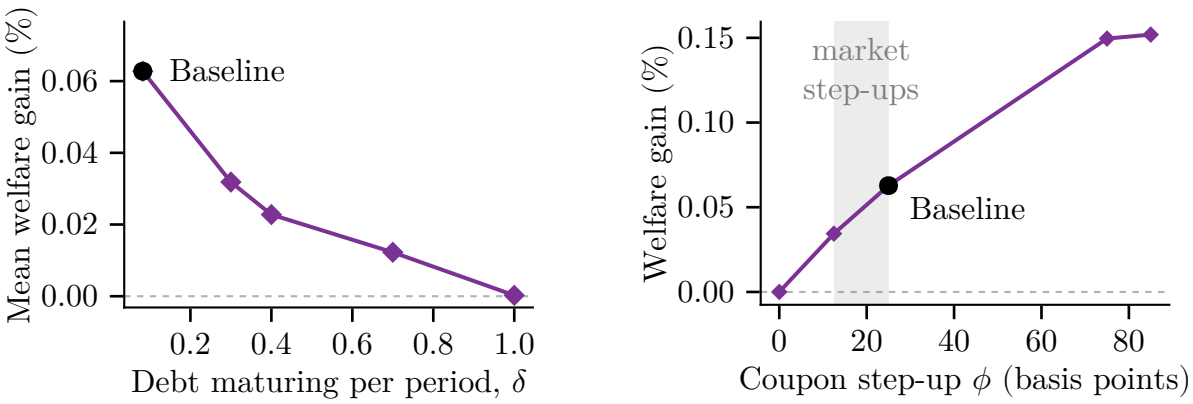


Note: Mean consumption-equivalent gain over the standard-bond economy from introducing the green bond with no green legacy, against debt at mean income, averaged over the disaster and no-disaster states.

Taking a stock Two results summarize the section. The government's use of the instrument is governed by the trade-off in (30): it earns the premium on the green share it carries forward and hands its future selves the step-ups behind that premium, so it holds green bonds where the premium is worth most to it and the burden least, at low income and light debt, and not where the covenant would do the most good. Even so, in the Markov equilibrium the government holds green bonds for about two thirds of its debt, a lower bound given the tie-breaking rule of Appendix D, the covenant moves investment toward the committed rule in the states where it binds, and the instrument raises welfare. Finally, the model-implied green shares are well above observed shares in the data where these types of bonds are rarely used.

7 Sensitivity Analysis

This final section varies two key parameters: the fraction of debt that matures each period, δ, and the coupon step-up, ϕ. For each parameterization, we re-solve the model with green bonds (holding all other parameters fixed) and compute the average welfare gain from introducing green bonds in that economy. Figure 10 presents the results.



(a) By maturity

(b) By covenant strength

Figure 10: The Value of the Green Bond

Note: Mean consumption-equivalent gain of the two-bond economy with a step-up over the same economy with no step-up, on the baseline ergodic path. Panel (a): $\phi = 25$ basis points at each maturity, each maturity re-solving the commitment rule, the two-bond economy and its no-step-up reference; the baseline $\delta = 0.0824$ is marked in black. Panel (b): each step-up at the baseline maturity; the calibrated $\phi = 25$ basis points is marked in black, and the shaded band spans 12.5 to 25 basis points.

Maturity. We re-solve the model at different values of δ. Three results summarize the sweep. First, the mean spread falls from 2.02 percent at the baseline maturity to 0.12 percent under one-period debt. Second, adaptation barely moves: the mean stock stays between 0.92 and 0.93, and the welfare gain from access to adaptation stays between 1.20 and 1.30 percent of consumption. Third, the value of commitment on the investment margin vanishes as debt shortens. We measure it with the rule economy of Appendix F, which commits the government to its investment rule and leaves borrowing discretionary: it is worth 0.050 percent of consumption at the baseline, 0.014, 0.010 and 0.006 at the middle maturities, and 0.000 under one-period debt, where the rule economy and the Markov equilibrium coincide. The value of the green bond declines in parallel: panel (a) of Figure 10 plots the mean gain of the $\phi = 25$ basis point instrument at each maturity, from 0.063 percent at the baseline through 0.032, 0.023 and 0.012 to 0.000 percent under one-period debt. The GEE of Section 4.2 explains these results: the investment wedge $H_{a'}$ is scaled by $(1 - \delta)$, and with one-period debt there is no outstanding stock to revalue, so the Markov investment choice is efficient.

Covenant strength. The second experiment varies the penalty the government pays when it skips the adaptation rule. We re-solve the portfolio economy at step-ups from zero to the 85 basis points that enforce the rule in every state, holding the baseline maturity (panel (b)

of Figure 10). The value of the instrument rises with ϕ throughout: from zero at $\phi = 0$, through 0.063 percent of consumption at the market's 25 basis points, to 0.152 percent at full enforcement. Adoption moves the other way: the green share of the debt carried forward peaks near the market's range, at 12.5 basis points and about three quarters of the stock, and falls to about a quarter at the strongest covenant. The direction follows from the trade-off in (30): a larger step-up scales up the premium and the burden together, but it also lowers the prices of the government's own net issues by more, because a heavier green legacy raises default risk by more, and that price effect is what turns adoption down. Adoption need not vanish even at a covenant that is never violated on the equilibrium path, because the covenant is respected only while green debt is outstanding: an all-standard stock would remove the penalty and with it the discipline. Stronger covenants are held less and are worth more: the market's step-ups sit where adoption peaks, not where welfare peaks.

8 Conclusion

This paper studies how sovereign default risk shapes government investment in climate adaptation. We measure adaptation from government budgets and find that past adaptation spending lowers the output loss a drought causes: a drought costs growth where spending has been low and not where it has been high. Sovereign spreads follow the same pattern. After catastrophic disasters and after drought onsets, spreads rise where past adaptation spending was low and not where it was high, in the budget panel and in an independent sample of severe cyclones. A default model with a durable adaptation stock, disciplined by the spending data, delivers both patterns: protection lowers losses and default risk, while the same fiscal state that moves spreads also moves adaptation, which is why annual regressions of spreads on adaptation detect little in the data and in the model alike. It implies that access to adaptation is worth 1.20 percent of consumption, with the gains concentrated in fiscally weak states facing active disasters.

On commitment, the model delivers a benchmark and a design result. Committing to the investment rule, with borrowing left discretionary, is worth 0.050 percent of consumption; this is the value of closing the investment wedge alone, not an upper bound on what a contract written on adaptation can deliver. The green bonds that markets write, coupons that step up when investment falls below a committed rule, are worth more than that benchmark, 0.063 percent of consumption at the market's step-up and 0.152 percent at full enforcement, because a coupon priced into the debt acts on the borrowing margin as well: the government holds them voluntarily, where the premium lenders pay for the step-ups is worth most to it, and they move investment toward the rule. Stronger covenants are worth more and are held

less.

Sovereign markets already write the contract we price, but point it elsewhere: the sustainability-linked bonds of Chile and Uruguay condition their coupons on emissions targets, disaster clauses condition on catastrophes, and the first operation conditioned on a verified adaptation outcome, a 2024 Barbados loan tied to the output of a water-reclamation plant, touches neither the issuance margin nor creditor protection. Two extensions follow from our results: a calibration matched to a single exposed economy, and contracts that condition on a verified adaptation rule, which no existing instrument offers.

References

Aguiar, M. and Amador, M. (2021). *The economics of sovereign debt and default*. Princeton University Press.

Aguiar, M., Amador, M., and Gopinath, G. (2009). Investment cycles and sovereign debt overhang. *The Review of economic studies*, 76(1):1–31.

Aguiar, M. and Gopinath, G. (2006). Defaultable debt, interest rates and the current account. *Journal of International Economics*, 69(1):64–83.

Arellano, C. (2008). Default Risk and Income Fluctuations in Emerging Economies. *American Economic Review*, 98(3):690–712.

Arellano, C., Bai, Y., and Mihalache, G. (2023a). Deadly debt crises: COVID-19 in emerging markets. *Review of Economic Studies*, 91(3):1243–1290.

Arellano, C., Mateos-Planas, X., and Ríos-Rull, J.-V. (2023b). Partial default. *Journal of Political Economy*, 131(6):000–000.

Arellano, C. and Ramanarayanan, A. (2012). Default and the maturity structure in sovereign bonds. *Journal of Political Economy*, 120(2):187–232.

Athey, S. and Imbens, G. W. (2017). The econometrics of randomized experiments. In Banerjee, A. V. and Duflo, E., editors, *Handbook of Field Experiments*, volume 1, pages 73–140. Elsevier, Amsterdam.

Barucci, E., Marazzina, D., and Nassigh, A. (2025). Sovereign debt default and climate risk. *arXiv preprint arXiv:2501.11552*.

Beirne, J., Renzhi, N., and Volz, U. (2021). Feeling the heat: Climate risks and the cost of sovereign borrowing. *International Review of Economics and Finance*, 76:920–936.

Boitan, I. A. and Marchewka-Bartkowiak, K. (2022). Climate change and the pricing of sovereign debt: Insights from european markets. *Research in International Business and Finance*, 62:101685.

Cavallo, E., Galiani, S., Noy, I., and Pantano, J. (2013). Catastrophic natural disasters and economic growth. *The Review of Economics and Statistics*, 95(5):1549–1561.

Cevik, S. and Jalles, J. T. (2022). This changes everything: Climate shocks and sovereign bonds∗. *Energy Economics*, 107:105856.

Chatterjee, S. and Eyigungor, B. (2012). Maturity, indebtedness, and default risk. *American Economic Review*, 102(6):2674–2699.

Conley, T. G. and Taber, C. R. (2011). Inference with "difference in differences" with a small number of policy changes. *The Review of Economics and Statistics*, 93(1):113–125.

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

Dvorkin, M., Sánchez, J. M., Sapriza, H., and Yurdagul, E. (2021). Sovereign debt restructurings. *American Economic Journal: Macroeconomics*, 13(2):26–77.

Eaton, J. and Gersovitz, M. (1981). Debt with Potential Repudiation: Theoretical and Empirical Analysis. *Review of Economic Studies*, 48(2):289–309.

Esquivel, C. (2024). Underinvestment and capital misallocation under sovereign risk. *Journal of International Economics*, 151:103973.

Esquivel, C. and Samano, A. (2023). Expansionary fiscal consolidation under sovereign risk.

Fisher, R. A. (1935). *The Design of Experiments*. Oliver and Boyd, Edinburgh.

Gómez González, J. E., Uribe, J., and Valencia, O. M. (2024). Asymmetric sovereign risk: Implications for climate change preparation. Technical report, IDB Working Paper Series.

Grosse-Steffen, C., Pagenhardt, L., and Rieth, M. (2025). Committed to flexible fiscal rules. *Journal of Monetary Economics*.

Hatchondo, J. C. and Martinez, L. (2009). Long-duration bonds and sovereign defaults. *Journal of international Economics*, 79(1):117–125.

Hatchondo, J. C., Martinez, L., and Sosa-Padilla, C. (2016). Debt dilution and sovereign default risk. *Journal of Political Economy*, 124(5):1383–1422.

Hsiang, S. M. and Narita, D. (2012). Adaptation to cyclone risk: Evidence from the global cross-section. *Climate Change Economics*, 3(02):1250011.

Kling, G., Lo, Y. C., Murinde, V., and Volz, U. (2025). Climate vulnerability and the cost of debt. *Oxford Open Economics*, 4.

Klomp, J. (2017). Flooded with debt. *Journal of International Money and Finance*, 73:93–103.

Kose, M. A., Kurlat, S., Ohnsorge, F., and Sugawara, N. (2022). A cross-country database of fiscal space. *Journal of International Money and Finance*, 128:102682.

Mallucci, E. (2020). Natural disasters, climate change, and sovereign risk. *FRB International Finance Discussion Paper*, (1291).

Mateos-Planas, X. and Ríos-Rull, J.-V. (2015). The generalized euler equation and the bankruptcy–sovereign default problem. *URL http://www. sas. upenn. edu/˜ vr0j/slides/vvucl. pdf. Slides, Department of Economics, University of Pennsylvania*.

Muff, S. and Keller, L. F. (2015). Reverse attenuation in interaction terms due to covariate measurement error. *Biometrical Journal*, 57(6).

Phan, T. and Schwartzman, F. F. (2023). Climate defaults and financial adaptation.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–322.

Vicente-Serrano, S. M., Beguería, S., and López-Moreno, J. I. (2010). A multiscalar drought index sensitive to global warming: The standardized precipitation evapotranspiration index. *Journal of Climate*, 23(7):1696–1718.

Young, A. (2019). Channeling Fisher: Randomization tests and the statistical insignificance of seemingly significant experimental results. *The Quarterly Journal of Economics*, 134(2):557–598.

A Alternative Measures of Adaptation

The main text measures adaptation with the budget panel alone. This appendix describes the two alternative measures used in the literature, documents how they relate to the budget measure, and repeats the main text's estimates with each of them.

The measures. The IMF climate expenditure database reports government spending on environmental protection as a share of GDP for a broad country panel, at the level of aggregate functional categories: the category is dominated by waste management, wastewater, and pollution abatement, while most adaptation spending is booked under other functions, so it is a broad environmental aggregate rather than a measure of adaptation. The readiness component of the ND-GAIN index summarizes a country's institutional capacity to deploy adaptation resources; it covers the longest panel but measures capacity rather than spending. The budget measure is the only one of the three built from the spending lines of adaptation-mandated agencies; the IMF measure is a functional aggregate that mostly records other things, and readiness is a capacity index. Appendix C documents how weakly the three co-move within countries.

What the budget measure contains. The numerator of the budget share is the executed spending of the agencies whose mandate is adaptation, identified from each country's administrative classification, not spending classified at the program or project level. Environment ministries and their technical services account for 73.6 percent of it, and their budgets also fund regulation, conservation, and administration. Three things follow. A constant non-adaptation fraction of agency spending is absorbed by the country fixed effects of the regressions and is invisible to the two moments the calibration targets, the coefficient of variation and the autocorrelation of the share, which are scale-free; it affects the measured level, which the calibration treats as a bound rather than a target, and the cross-country ranking that the event split uses. To see how much rests on the environment ministries, we recompute the measure from the water and civil-protection agencies alone, 26 percent of identified spending. It exists for 16 of the 19 countries, with a median share of 0.12 percent of spending, and its two calibration moments are close to those of the full measure: a lag-1 autocorrelation of 0.65 against 0.66, and a within-country coefficient of variation of 40 percent against 32. The event split cannot be repeated on this subset on equal terms: Brazil and Poland, the two low-share events with the largest spread responses, identify adaptation spending only in their environment agencies and drop out, and on the six remaining events the rank correlation is -0.26 with $p = 0.33$. The regression of equation (31) is imprecise with

either measure.

Data construction. Three conventions in the underlying series should be recorded. Annual sovereign spreads are unweighted means of monthly means, and a country-year is retained only if at least six months are covered; the underlying daily EMBI series carries previous-value fills on weekdays without a quote. The ND-GAIN readiness aggregate is imputed by its provider where component pillars are missing. IMF environmental-expenditure figures are reported irregularly, and we carry the most recent observation forward for up to five years. None of these fills touches the outcome or the regressors of the drought design of Section 2.2 or of the damage regressions of this appendix, whose samples are restricted to country-years with both a recorded disaster and an observed budget share.

Figure 13 compares the panels. The budget panel covers 229 country-years across 19 economies between 2003 and 2024; the IMF aggregates cover 1,928 country-years across 128 countries from 1995. The two record different objects in different units, a budget share of identified adaptation lines against a GDP share of aggregate environmental expenditure, and they behave differently: the median budget adaptation share roughly doubles over the sample, from about half a percent of the budget in the mid-2000s to close to one percent by 2021, while the median IMF measure is flat near half a percent of GDP for twenty-five years. Across the sixteen countries where both exist, country means are essentially uncorrelated. The rise in budgeted adaptation alongside a flat aggregate is consistent with reclassification toward identifiable adaptation lines as much as with growth in underlying effort, which is one more reason the calibration targets the level and volatility of the best-measured series rather than any trend.

Missing data. The adaptation series themselves are never filled: no budget share is interpolated, carried forward, or winsorized, and every merge drops unmatched rows rather than imputing them. Five conventions elsewhere are worth stating. EM-DAT records a start month but not always a start day; for the 11 percent of climate events where the day is missing we date the event to mid-month, which matters only for the alignment of the daily spread window. In the cyclone design of Appendix B, pre-event IMF spending is the most recent observation in the five years before landfall, which binds for six of the fifteen events. An absent EM-DAT record is read as no disaster, so the disaster indicator is zero rather than missing. Where the IMF reports environmental-protection components separately, a missing component is treated as zero rather than dropping the country-year. Finally, ND-GAIN publishes a complete readiness aggregate even where its own underlying pillars are missing; we use the published aggregate as is.

Exposure and spending in the broad panels. Across the 127 countries with both IMF spending and disaster data, the most exposed tercile spends 0.30 percent of GDP against 0.64 percent for the least exposed, and the rank correlation between spending and disaster frequency is -0.28; controlling for income leaves the partial correlation at -0.21. Within the emerging markets that carry EMBI spreads, the correlation turns positive, at 0.54 in levels: adaptation is an equilibrium response to risk where finance permits it, and follows fiscal capacity globally. This is the broad-panel counterpart of the budget-panel fact in the main text.

Adaptation capacity is priced in spread levels. Conditional on the debt-to-GDP ratio and income, sovereign spreads decrease with the ND-GAIN readiness index; the relationship is negative in all twelve specifications spanning fixed-effects structures, samples, and spread sources, and statistically significant in eight. With country and year fixed effects, the coefficient on readiness in log spreads is -0.026 with a t-statistic of 4.2. Markets treat adaptation capacity as relevant for default risk before any disaster occurs.

Recorded losses and the budget share. Before the drought design of Section 2.2, we related recorded disaster losses to the budget share directly:

$$\text{Damages}_{it} = \omega_0 + \omega_B \, \Delta \, \text{IA}_{it} + \omega_S \, \text{Size}_{it} + \alpha_i + \epsilon_{it}, \tag{31}$$

where Damages_{it} are logarithmic losses, measured in percent of GDP. ΔIA_{it} is the change within the country in the share of the adaptation budget, expressed as a fraction of government spending and dated the year before the disaster so that it is predetermined; Size_{it} is the affected population in percent of population; and α_i is a country fixed effect. A coefficient of $\omega_B = -6.99$ therefore means that a one-percentage point increase in the adaptation share reduces the damage by about 0.07. Table 5 presents the estimates, with the sample restricted to country-years with a weather disaster.

Table 5: Cost of Natural Disasters and Adaptation Spending

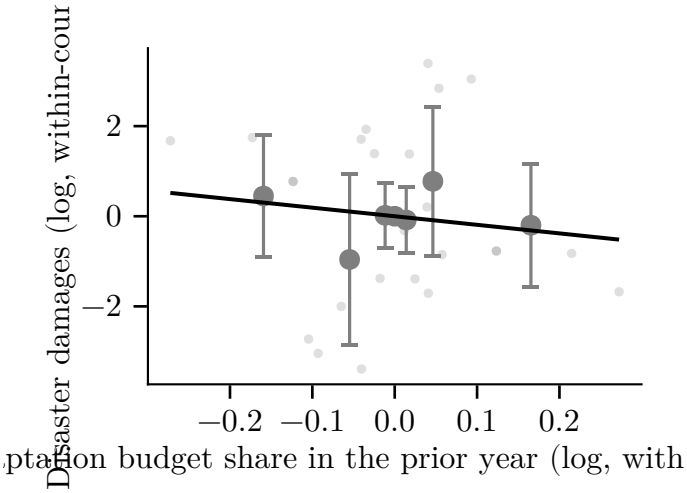
	ω_B	ω_S
Coefficient	−6.99	0.007
	(0.28)	(0.42)
Observations	48	

Notes: P-values in parentheses. The p-value for ω_B is one-sided (against $\omega_B < 0$), from a wild cluster bootstrap (Rademacher weights by country, null imposed, 4,999 replications); the p-value for ω_S is a two-sided p-value from the cluster-robust asymptotic t, a different inference reported for reference only. The specification includes country fixed effects; damages enter as $\log(1 + \cdot)$ and the sample is restricted to country-years with a weather disaster in 14 budget-panel countries. The change in the adaptation share is dated the year before the disaster, so it is predetermined.

The coefficient has the sign protection implies but is imprecise: a one-sided wild cluster bootstrap p-value of 0.28 on 48 disaster country-years. It is not used to calibrate the model; Section 5.1 reports the model's counterpart as an untargeted check. Two features of the data limit what this regression can show and motivate the design of the main text. A country-year enters only when a disaster is recorded, which selects on the outcome that protection reduces, and the damage figure is missing for two recorded events in three.

The damage scatter in the budget data. Figure 11 is a descriptive companion to the estimates of equation (31), not a plot of their variation: it shows log damages against the prior-year log adaptation share, both demeaned by country and year, across disaster country-years, whereas regression (31) uses the change in the share, country fixed effects only, and the size control. The fitted slope is −1.9; the dispersion around it is wide, which is why the text works with the regression estimates and their bootstrap inference rather than with this figure.

Figure 11: Adaptation Spending and the Cost of Disasters: Within-Country Variation



Note: Each point is a country-year with a climate disaster in the nine budget-panel countries with disaster damage data, 2004–2023. Vertical axis: log direct economic losses as a share of GDP from EM-DAT. Horizontal axis: log adaptation share of government spending in the year before the disaster. Both variables are demeaned by country and year; the line is the fitted slope, −1.9.

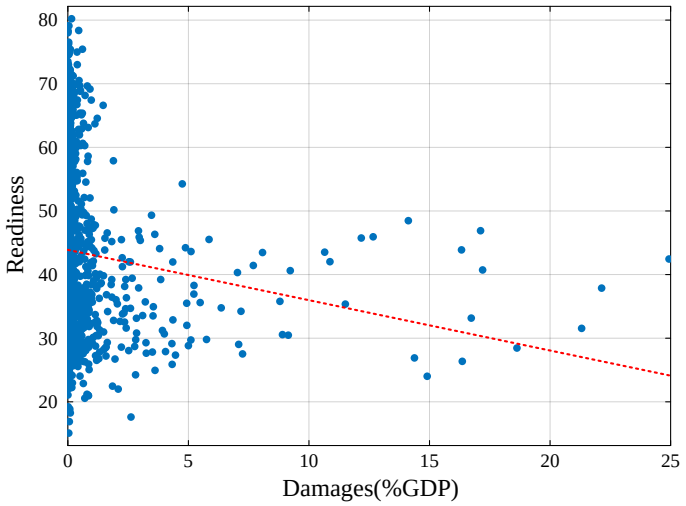
The damage interaction across measures. A second way to ask whether adaptation protects is to let it moderate the damage-intensity gradient rather than shift the level of damages. For each measure A we estimate

$$\log(1+\text{Damages}_{it}) = \omega_1 I_{it} + \omega_2 I_{it} \times A_{i,t-1} + \omega_3 A_{i,t-1} + \omega_4 \log \text{GDPpc}_{it} + \omega_5 \log \text{Pop}_{it} + \alpha_i + \lambda_t + \epsilon_{it} \quad (32)$$

on the country-years with a weather disaster, where I_{it}, the log of one plus the affected population per capita, is the physical exposure of the event, $A_{i,t-1}$ is the lagged adaptation measure standardized to unit variance, and α_i and λ_t are country and year effects; standard errors are clustered by country. Adaptation that protects flattens the gradient, $\omega_2 < 0$. The estimates of ω_2 are -0.052 with the budget shares, -0.050 with the IMF aggregates, and 0.009 with readiness. The three are not on a common scale: the budget-share coefficient is per within-country standard deviation of the share, while the IMF and readiness coefficients are per pooled standard deviation, so their magnitudes are not directly comparable and none is estimated precisely enough to rank against the others. What the comparison does show is that the two spending measures carry the sign protection implies while the capacity index does

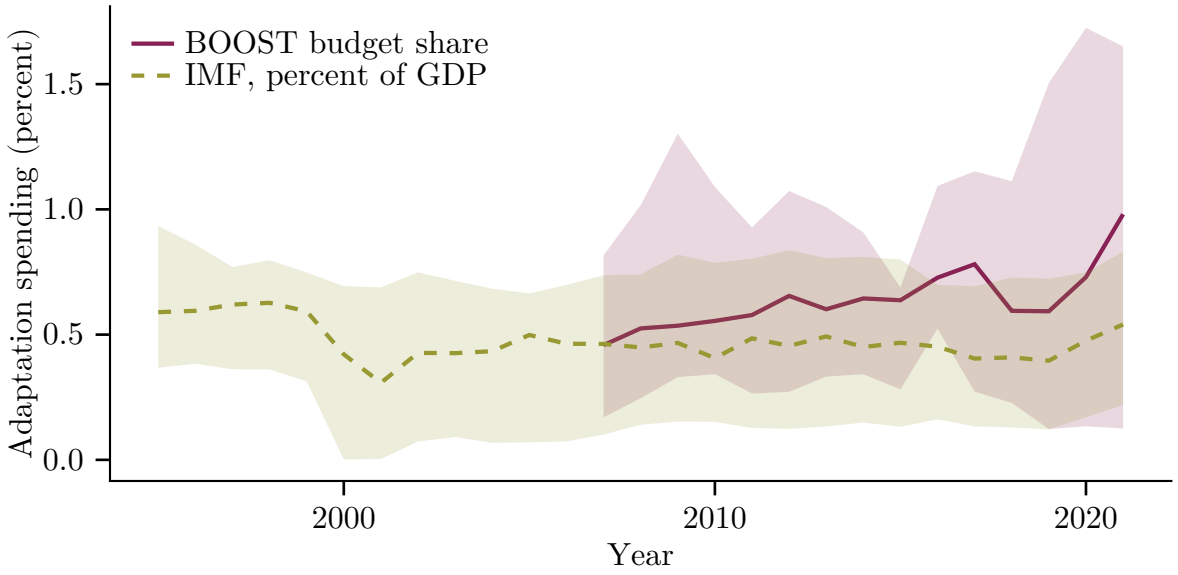
not, which is what a moderator with no within-country information about spending should deliver (Muff and Keller, 2015). Figure 12 shows the readiness version of the unconditional damage pattern, the exhibit that a capacity index supports: across countries, disasters cause smaller losses where readiness is higher, a relationship that mixes protection with income and institutions and is reported here for comparability with the prior literature.

Figure 12: Damage and Vulnerability: the Readiness Version

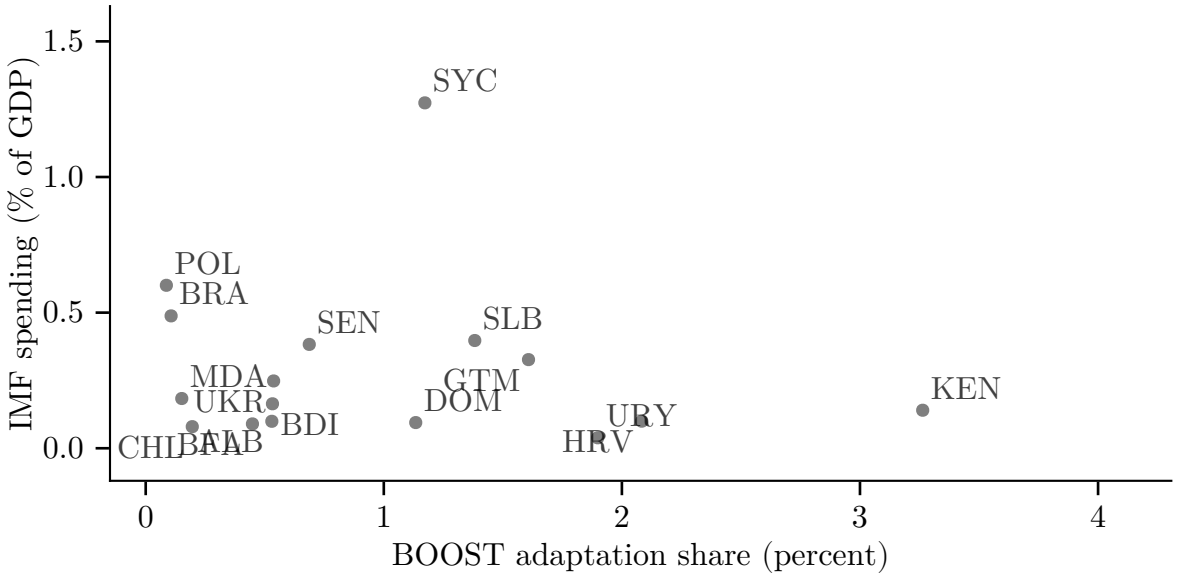


Note: Each point is a disaster in one of 36 emerging markets between 1995 and 2023. Damages are direct economic losses as a share of GDP from EM-DAT; readiness is the ND-GAIN index in the year of the event. The line is a linear fit.

Figure 13: Adaptation Spending in the Data



(a) Evolution



(b) Country means, both measures

Note: Panel (a): median across countries by year with interquartile bands; BOOST is the identified adaptation share of the government budget, the IMF measure is environmental-protection expenditure as a share of GDP; years with fewer than eight reporting countries are omitted, as is the final BOOST year, in which the reported budgets are incomplete. Panel (b): country means over the sample for the sixteen countries covered by both measures, each point tagged with its ISO3 code; tags for countries whose means nearly coincide overlap.

B The Cyclone Event Study and the Event Sample

This appendix contains the second event study referenced in Section 2.3: tropical cyclones selected on physical severity and split by an independent spending measure, an event set fully disjoint from the main text's budget-panel sample. It then documents how the event sample changes with the definition of a catastrophic disaster and the severity threshold, reports the robustness checks for both designs, and records the design history of the budget-panel protocol.

The cyclone design. Selecting on landfall wind speed removes any margin through which adaptation could affect selection: the sample is tropical cyclones making landfall with maximum sustained winds of at least 125 knots in a country with daily EMBI coverage. This yields 15 events between 2006 and 2025, ten in the Philippines, four in China, and one in Colombia. Events are split by the country's IMF-measured adaptation spending in the years before the event, relative to the sample median; the budget panel does not cover the cyclone belt, which is why this design uses the aggregate measure. Spread paths are constructed exactly as in the main text.

The result is as follows. Before the event, the two groups are indistinguishable: the difference in the pre-event month is 1.0 log points with a standard error of 0.9. Because such pre-tests are typically underpowered (Roth, 2022), we do not treat this as validation by itself; the sharper statement is a breakdown calculation: a linear differential trend would have to be roughly four times larger than anything the pre-event window admits to account for the month-two difference. After the event, spreads rise only where adaptation investment was low. In the second month after landfall, spreads in low-investment countries are 3.8 log points above their pre-event level, with a confidence interval that excludes zero, while spreads in high-investment countries remain at their pre-event level. The difference is -4.5 log points. Because 15 events cannot support asymptotic inference, we compute the exact randomization distribution over all 5,005 possible assignments of events to groups: the one-sided p-value for the month-two difference is 0.050, and 0.058 for the first post-event quarter.

Three checks discipline the interpretation. First, the contrast is concentrated in the most intense storms. Table 6 repeats the exercise at wind thresholds from 90 to 137 knots; the rows are nested, each adding the storms above its cutoff, so they test robustness to the cutoff rather than separate severity bands. The difference is present at every cutoff from 113 knots up, largest and most precise between 120 and 130 knots, and it vanishes once the sample is extended to moderate storms, which more than doubles it. If the difference reflected some generic fragility of low-adaptation countries to any shock, adding moderate storms would

not wash it out. Second, the country composition of the sample raises a clustering concern: with ten Philippine and four Chinese events, treatment is effectively assigned at close to the country level, and inference that permutes events freely partly tests differences between countries (Conley and Taber, 2011). The within-country design is the answer to this concern, not merely a robustness check. Restricting to Philippine events alone and splitting by the country's own spending over time removes every cross-country difference by construction, and yields, across the eleven Philippine landfalls at or above 120 knots, a first-quarter difference of -5.0 log points with an exact one-sided p-value of 0.073 (at the 125-knot baseline, where only ten Philippine events remain and one falls in the high group, the difference is -3.0 log points with $p = 0.200$), where the small number of admissible assignments places a lower bound of 0.018 on any attainable p-value. Third, the difference does not depend on any single event: dropping events one at a time leaves every estimate negative, and the groups face storms of balanced physical severity, 143 against 139 knots on average. Deaths per million are three times higher in the low-investment group, which is the protection mechanism itself rather than a confound.

Construction details. Five conventions are worth recording. Winds are the maximum sustained wind recorded within 100 kilometres of the country's coastline around first approach, taking the larger of the WMO and US agency values in IBTrACS; this is a proxy for landfall intensity rather than a landfall reading, and 125 knots is approximately the 90th percentile of landfall winds in the spread sample. Spread paths are the log spread minus the country's own mean over trading days -10 to -3, adjusted in calendar time against countries with no hurricane-force landfall in the same window, and binned into 21-trading-day months. Where EM-DAT records severity only at the country-year level the event is dated to mid-year; the main text's eight-event sample uses event-level dates throughout, and the paragraph below records why. IMF spending is carried forward for up to five years when the series is missing in the pre-event window. At the median split, ties are assigned to the low group ($\leq$ the median). Randomization inference enumerates all assignments when the number of admissible splits is at most 200,000 and draws 20,000 seeded random assignments otherwise. The split and rank tests of the main text (70 splits and 40,320 orderings), the baseline cyclone p-values (125 knots and above), and the Philippine-only p-values are from full enumeration; the 90- to 113-knot rows of Table 6, with between 705,432 and 4.1×10^9 admissible splits, are Monte Carlo approximations with 20,000 draws, whose simulation standard error is below 0.004 at every reported value.

The budget-panel protocol: variants and design history. The main text's eight-event design was fixed before estimation, and its result is invariant to the discretionary margins. Replacing the lagged share by a depreciation-weighted stock of past shares (at the calibrated depreciation rate) or by a three-year average leaves the primary rank correlation, and hence its exact p-value, unchanged at -0.55 and 0.086; a five-year average gives -0.60 and $p \approx 0.07$. Removing the crisis-window exclusion adds one event, Ukraine 2008, whose window spans the Lehman default, and moves the primary to -0.48 with $p \approx 0.10$. One design lesson is worth recording: an earlier version of the all-disaster-type sample dated events to mid-year when event-level dates were unavailable, and that version produced a large, apparently significant difference that dissolved once events were dated from event-level EM-DAT records—the July convention had placed Poland's 2010 floods in their own pre-window and let Ukraine's 2008 crisis run-up masquerade as a pre-event level. Approximate dating is not a small implementation detail in daily event studies, and no result in this paper rests on it.

Disaster definitions. Pooling all climate disaster types and selecting country-years above the 90th percentile of deaths per million, the filter of Cavallo et al. (2013), yields 11 events in the spread sample: six in the Philippines, two in El Salvador, two in Peru, and one in Venezuela. Nine of the eleven are storms; broadening the definition beyond storms adds almost no events, because deadly climate catastrophes in countries with traded sovereign debt are overwhelmingly tropical cyclones. The deaths-based sample has two defects. Five of the eleven events carry approximate dates, because EM-DAT severity is recorded at the country-year level, and only two events have pre-event spending data, so the sample can only be split by the readiness index. That split produces a difference with the opposite sign: spreads rise more after catastrophic events in high-readiness countries, by 18.5 log points at the second month. This reversal is what selection on an outcome predicts. Deaths are reduced by adaptation, so a deaths filter retains the routine disasters of unprotected countries but only those disasters severe enough to overwhelm protection elsewhere, and the treated events of protected countries are systematically harsher. Selecting on physical severity removes this margin. At the 99th percentile the deaths filter leaves a single event, the 1999 Vargas disaster in Venezuela.

Severity thresholds. Table 6 reports the physically filtered sample as the landfall wind threshold varies, together with the month-two high-minus-low difference and its one-sided randomization p-value. Two features stand out. The sample composition is stable: at every threshold the sample is dominated by the Philippines and China, because these are the countries with traded sovereign debt where intense cyclones make landfall; lowering

the threshold to 90 knots doubles the sample but adds only Costa Rica. The difference is concentrated in intense storms: it is present at every cutoff from 113 knots up and absent once the sample extends below Category 4 intensity; because the rows are nested, this is robustness to the cutoff, not a comparison of separate severity bands. No event enters the sample while an earlier event in the same country has an active window, at any threshold.

Table 6: Event Sample and Post-Event Spread Difference, by Wind Threshold

Threshold (kt)	Events	Low/High	PHL	CHN	Diff. month 2	p (one-sided)
90	35	19/16	13	20	−0.3	0.45
96	31	17/14	12	17	−0.1	0.48
105	24	12/12	11	12	−0.5	0.45
113	22	11/11	11	10	−4.4	0.16
120	17	9/8	11	5	−8.6	0.01
125	15	9/6	10	4	−4.5	0.05
130	14	9/5	10	3	−4.6	0.05
137	9	7/2	8	1	−5.5	0.11

Note: Tropical cyclones with landfall winds at or above each threshold in countries with daily EMBI coverage and pre-event IMF spending data. Differences are in log points times 100. The p-value is computed by enumerating all assignments of events to groups when their number is at most 200,000 (the rows at 120 knots and above) and from 20,000 seeded random draws otherwise (the rows from 90 to 113 knots), so the latter are Monte Carlo approximations with a simulation standard error below 0.004. Remaining events at each threshold are in Colombia (one at 90 through 130 knots) and Costa Rica (one at 90 and 96 knots).

Baseline checks. At the baseline threshold of 125 knots, dropping events one at a time leaves the month-two difference between −3.0 and −5.7 log points, negative in all 15 cases. The two groups face storms of balanced physical severity: mean landfall winds are 143 knots in the low-investment group and 139 in the high-investment group, with a permutation p-value of 0.62. Damage as a share of GDP is, if anything, larger in the high-investment group, 0.22 against 0.13 percent of GDP. Deaths per million are three times larger in the low-investment group, 3.0 against 1.1, which reflects the protection margin itself. Pre-event spread levels are higher in the high-investment group by 0.33 log points, an imbalance that works against finding a larger post-event rise in the low group, since spreads are more volatile at higher levels.

The drought index. The drought index of Section 2.1 comes from the global SPEI12 grid, monthly at half a degree, 1901–2024. Grid cells are assigned to a country when their centers fall inside its polygon and weighted by area; the country's monthly value is the weighted mean over its cells, and the drought index of a year is the worst monthly value with the sign

flipped. We pre-specified four measures: this continuous index, the shares of the country's area in severe drought (SPEI ≤ -1.5) and in severe excess (≥ 1.5) in December, and the share in either. Only the continuous index carries the signal reported in the text; the December shares exceed one half in only 20 country-years and are too coarse to identify an interaction, and wet extremes raise growth in these economies. A placebo with geophysical disasters is uninformative, with 16 country-years hit.

Robustness of the growth design. Table 2 reports p-values from standard errors clustered by country. With 19 clusters we also compute wild cluster bootstrap p-values (Rademacher weights by country, null imposed, 4,999 replications) for the interaction in its four columns: 0.052, 0.041, 0.007 and 0.126 one-sided, 0.106, 0.085, 0.014 and 0.244 two-sided; the delete-one-country jackknife gives 0.097, 0.089, 0.007 and 0.296. Dropping any one of the 19 countries leaves the interaction between 0.39 and 0.67 with cluster t-statistics between 1.3 and 2.4; over two years of growth, between 1.09 and 1.58 with t-statistics between 2.7 and 4.7. Splitting the panel at the median of lagged income, the interaction is 1.04 ($t = 2.8$) in the poorer half and 0.02 ($t = 0.1$) in the richer half. Adding the previous year's drought index as a control gives 0.82 ($t = 2.5$). In the units of the budget, past spending one percentage point of government spending higher lowers the growth cost of a unit of drought by 0.96 percentage points ($t = 1.6$). Without an interaction, the slope of growth on the drought index is -0.76 percentage points in country-years where past spending is below the country's median and 0.30 where it is above; split into terciles, the slopes are -0.94, 1.22 and -0.49, so the pattern is not monotone across three groups. It is present in the poorer half of the panel and not in the richer half.

Table 7 varies how past spending is aggregated. The five-year mean is the baseline because it is the simplest measure with a fixed window. The interaction is weaker with a three-year window, similar with the previous year's share alone, and stronger with every measure that reaches further back, including the depreciation-weighted stock of the model at the calibrated rate $\delta_a = 0.75$ and at slower rates. The table lists the weight each measure places on the budget share of each past year.

Table 7: The Drought Interaction under Alternative Measures of Past Spending

	Weight on the share of year								
Measure	$t-1$	$t-2$	$t-3$	$t-4$	$t-5$	older	β_2	p-value	wild p
Mean of the five previous shares (baseline)	0.20	0.20	0.20	0.20	0.20	0	0.58	0.064	0.052
Mean of the three previous shares	0.33	0.33	0.33	0	0	0	0.37	0.232	0.131
Previous year's share	1	0	0	0	0	0	0.54	0.029	0.014
Mean of all previous shares	equal weights on every previous year						0.78	0.010	0.012
Sum of all previous shares	weight one on every previous year						0.76	0.003	0.006
Stock, $\delta_a = 0.75$	0.75	0.19	0.05	0.01	0.00	0.00	0.77	0.010	0.004
Stock, $\delta_a = 0.50$	0.50	0.25	0.12	0.06	0.03	0.03	0.98	0.002	0.001
Stock, $\delta_a = 0.25$	0.25	0.19	0.14	0.11	0.08	0.24	1.15	< 0.001	< 0.001

Note: Each row re-estimates column (1) of Table 2 with the adaptation measure named, built from budget shares up to $t-1$ and standardized within country. Weights are the measure's weights on the share of each past year, normalized to sum to one; the stock weights year $t-k$ by $(1-\delta_a)^{k-1}$, with the tail beyond five years in the column "older". β_2 is the coefficient on drought $\times$ measure. The p-value is two-sided from standard errors clustered by country; wild p is the one-sided wild cluster bootstrap p-value. 227 country-years, 19 countries, 2004–2023.

The drought event study. Events are onsets: the first month in which a country's twelve-month SPEI is at or below -1 after twelve months above it; a later onset within eighteen months of a kept one is dropped. An onset is hard when the index reaches -1.5 within the following twelve months. Paths are monthly log spreads in log points relative to the event's own mean over months -6 to -1, net of the cross-country mean of the same construct over the EMBI countries with no onset within twelve months of the same calendar month, with at least five such controls. Three exclusion rules are applied before any result is computed: windows overlapping the 2008–2009 crisis or the March 2020 repricing, as in the main text; windows overlapping a war on the country's territory or at its border, Ukraine from February 2014 through 2015 and Ukraine and Poland from February 2022 through 2023; and windows overlapping a sovereign restructuring, which the war rule already covers for Ukraine in 2015. The war rule matters: without it, Ukraine's onsets of October 2014 and June 2022 and Poland's onset of May 2022 carry spread movements of 60 to 130 log points that have nothing to do with rainfall. Table 8 lists the hard onsets in the budget-panel countries with daily EMBI coverage.

Table 8: Hard Drought Onsets in the Budget-Panel Countries

Country	Onset	Peak SPEI	Spending	Group	Months 1–3	4–6	7–12
Brazil	2015-11	−1.51	0.13	low	25.6	12.6	2.1
Dominican Rep.	2018-11	−2.20	0.78	high	−1.9	−3.9	−5.9
Dominican Rep.	2021-11	−1.63	1.09	high	−9.1	−1.8	−14.7
Mexico	2011-07	−1.63	0.55	high	3.0	−2.4	−4.5
Mexico	2023-09	−1.56	0.38	low	9.8	17.0	20.4
Poland	2005-10	−1.74	0.05	low	33.5	50.5	44.9
Poland	2018-08	−2.20	0.11	low	−15.5	−9.7	−39.7
Ukraine	2007-05	−1.55	0.44	high	−5.3	7.3	13.8
Ukraine	2011-11	−1.63	0.39	low	30.8	35.0	37.6
Uruguay	2022-09	−2.13	1.95	high	−20.8	−23.6	−27.8

Note: Peak SPEI is the lowest monthly value of the twelve-month index in the year after onset. Spending is the mean budget share over the five years before the onset, in percent of government spending, the most recent available value when the year before the onset is missing; the group splits at the sample median. The last three columns are the event's mean adjusted spread path over the months shown, in log points.

Split at the median of past spending, the high-minus-low difference is −23.7 log points in months 1–3, −26.0 in months 4–6, and −20.9 in months 7–12, with exact one-sided randomization p-values of 0.028, 0.032, and 0.123 over the 252 splits; the rank correlations between past spending and the response are −0.59, −0.58, and −0.48 (randomization $p = 0.041$, 0.044, 0.084). Splitting instead on the measure standardized within country, the differences are −16.8, −26.2, and −39.1 ($p = 0.091$, 0.028, 0.012). Dropping any one of the ten events leaves the months 4–6 difference between −20 and −32 log points with p at most 0.063. Over all 17 onsets, hard or not, the difference is −15 log points in each of the first two quarters ($p = 0.035$ and 0.053). In the full EMBI panel, with the IMF measure as the only spending series available, drought onsets raise spreads by 8 log points in months 4–6 ($t = 2.2$) and 9 in months 7–12 ($t = 1.9$) across 26 onsets in 13 countries, and the IMF measure does not split the response, which is what its lack of within-country signal implies.

C The Readiness Index and Adaptation Spending

The published version of this paper used the ND-GAIN readiness index as its measure of adaptation. This appendix documents how the index relates to measured adaptation spending and states precisely the role the index retains.

The readiness index summarizes economic, governance, and social capacity to deploy adaptation resources. It is not a measure of spending, and the data bear this out. Between

countries, mean readiness is uncorrelated with the BOOST adaptation budget share among the 19 emerging markets where both exist, with a correlation of -0.11; it correlates at 0.35 with IMF-measured environmental spending across the broad panel of 126 countries, a relationship that reflects income and institutions rather than budget effort. Within countries over time, where the fixed-effects specifications and the event groupings operate, the index contains almost no information about spending: the within-country correlation is 0.09 with the BOOST share and 0.11 with the IMF measure. The two spending measures also agree only imperfectly with each other, with a within-country correlation of 0.16, consistent with substantial measurement error in both and with their different denominators and coverage.

These correlations fix the index's role in the revision. Readiness enters as a measure of capacity where capacity is the object of interest: it is the moderator in the damage scatter of Figure 12, and it is the variable that markets price in the spread-level regressions of Appendix A, where its slow-moving institutional content is a feature rather than a defect. Readiness is not used to calibrate the model, to group events, or to measure investment; those roles fall to the spending measures, with the budget panel preferred where available. The one exception is noted in Section 5.1: the bin frequencies of the discretized disaster-intensity grid are taken from the damage regression estimated with this index. The near-zero within-country correlation between readiness and spending also explains two patterns in the results: the damage interaction estimated with readiness is zero while the same specification with BOOST yields -0.052, which is what attenuation implies when the proxy carries no signal about the variable of interest; and the one readiness-based event grouping we report, the deaths-selected sample of Appendix B, produces a difference of the wrong sign, which is what selection on an outcome combined with an uninformative proxy implies, and which is why no event design in the paper groups events by readiness.

D Algorithm

We solve the model using value function iteration on the discretized state space. Throughout this appendix the income state is the pair $(\hat{y}, \gamma)$ of the pre-disaster income node and the current disaster intensity ($\gamma = 0$ when no disaster is active), and post-disaster income is $y = \hat{y}\,[1 - h(a, \gamma)]$, which depends on the entering stock a; objects that the main text indexes by post-disaster income y are indexed by $(\hat{y}, \gamma, a)$ here, and y denotes post-disaster income wherever it appears in a budget. The state under repayment is $(b, a, \hat{y}, \gamma)$ and under default $(a, \hat{y}, \gamma)$; the adaptation stock is therefore carried in every value, policy and price object. We use $n_b = 20$ evenly spaced debt points on $[0, 0.7]$; $n_a = 40$ adaptation points on $[0, 4]$, spaced geometrically so that the depreciated stock $(1 - \delta_a)a$ lands on a grid node; $n_y = 10$ income levels from Tauchen's method over 2.5 standard deviations; and $n_\gamma = 4$ disaster intensities plus a no-disaster state. Debt and adaptation are discretized jointly: the choice grid has n_a blocks of n_b debt points, so a single index carries the pair (a', b').

1. Construct grids for $b' \in \mathcal{B}$, $a' \in \mathcal{A}$, and discretize the shock processes (ϵ, γ).

2. Guess initial value functions $V^R(b, a, \hat{y}, \gamma)$ and $V^D(a, \hat{y}, \gamma)$, and an initial price function $q(\hat{y}, \gamma, a, b', a')$. The price carries the entering stock a because the conditional distribution of next period's income depends on current post-disaster income $y = \hat{y}[1 - h(a, \gamma)]$, which the predetermined stock determines.

3. For each state $(b, a, \hat{y}, \gamma)$ under repayment, solve for optimal policies (b', a'):

$$V^R(b, a, \hat{y}, \gamma) = \max_{b', a'} \left\{ u(c^R) + \beta \mathbb{E}[V(\chi')] \right\}$$

   subject to $c^R = y - \kappa b - f(i) + q(\hat{y}, \gamma, a, b', a')(b' - (1 - \delta)b)$ and $i = a' - (1 - \delta_a)a \geq 0$, with $y = \hat{y}[1 - h(a, \gamma)]$ the income net of the current disaster loss evaluated at the entering stock. The pair (b', a') is chosen by grid search over the joint choice grid, and the joint discrete choice is smoothed with an extreme-value taste shock of scale ρ_b attached to each (b', a') pair, which delivers the choice probabilities used in the price update; the borrowing and investment choices are not smoothed separately.

4. For each state $(a, \hat{y}, \gamma)$ under default, solve for optimal adaptation. Because the current default cost ν does not affect the choice of a', we work with the ν-independent value of equation (14):

$$V^D(a, \hat{y}, \gamma) = \max_{a'} \left\{ u\left(y^D - f(i)\right) + \beta \mathbb{E}\left[\theta V\left(0, a', \hat{y}', \gamma'\right) + (1 - \theta) V^D\left(a', \hat{y}', \gamma'\right)\right] \right\}$$

with $i = a' - (1-\delta_a)a \geq 0$ and y^D the default endowment at post-disaster income y. The government re-enters markets with zero debt and the option to default again, so the continuation on re-entry is V, not V^R. The full default value is $V^D(a, \hat{y}, \gamma) - \nu$.

5. Compute the default threshold $\hat{\nu}(b, a, \hat{y}, \gamma) = V^D(a, \hat{y}, \gamma) - V^R(b, a, \hat{y}, \gamma)$ and the default probability $F(\hat{\nu})$. The default decision is smoothed with an extreme-value shock of scale ρ_d, the numerical counterpart of the stochastic default cost ν in Section 3: the implied distribution of the cost is logistic with scale ρ_d, $F(\hat{\nu}) = 1/(1 + e^{-\hat{\nu}/\rho_d})$.

6. Update the price function using the break-even condition:

$$q(\hat{y}, \gamma, a, b', a') = \mathbb{E}\left[\frac{(1 - d(\chi'))[\kappa + (1-\delta)q(\hat{y}', \gamma', a', b'', a'')]}{R}\right]$$

   where the expectation over $(\hat{y}', \gamma')$ is taken conditional on current post-disaster income $y = \hat{y}[1 - h(a, \gamma)]$, evaluated at the *entering* stock a, while the damage the stock a' mitigates enters next period's post-disaster income.

7. Update the value functions and the price and check convergence. Stop when both the value residual $\|V^{new} - V^{old}\|$ and the price residual $\|q^{new} - q^{old}\|$, measured on the undamped update, are below $\varepsilon = 10^{-5}$; otherwise return to step 3. The commitment economies of Appendix F reach the price criterion through the price polish described below.

Numerical details. Four features of the implementation are worth recording. First, the taste shock smooths the *joint* (b', a') choice rather than the borrowing choice alone, with scale $\rho_b = 5 \times 10^{-3}$; the default decision is smoothed the same way with scale $\rho_d = 10^{-4}$. Because a smoothed choice carries an inclusive-value term that is not part of the government's payoff, all welfare numbers in the paper are computed from a separate evaluation step that holds the optimal policies fixed and strips that term, so the taste shock affects the equilibrium the government faces but not the values we compare.

Second, the price schedule carries the inherited stock as a separate argument, $q(\hat{y}, \gamma, a, b', a')$. In the notation of Section 3 this is still the $q(y, b', a')$ of equation (18), because the y of the main text is income *after* the current disaster; the numerical state instead indexes income before the loss together with the current intensity, so the inherited stock has to appear explicitly. This matters because the conditional distribution of next period's income depends on current post-disaster income, $\hat{y}[1 - h(a, \gamma)]$, which is determined by the *predetermined* stock a, while the stock a' chosen today governs the damage next period. Evaluating post-damage income

at a' instead would let the current disaster be undone by tomorrow's protection: at the most severe intensity the conditioning income would move by 12.9 percent of output, about two nodes of the income grid, between the lowest and highest a', and the price schedule would inherit a dependence on a' that is an artifact of the timing rather than the protection channel of Lemma 1. Carrying a explicitly adds one dimension to the price array and removes that artifact.

Third, continuation income falls between the discretized income levels once damages are subtracted, and we interpolate linearly between the two adjacent rows of the transition matrix rather than rounding to the nearest node.

Fourth, a word on convergence. Residuals are measured on the *undamped* update, that is on $\|T(q) - q\|$. This matters because an earlier version of the solver damped the price step and measured the residual afterwards, which reports the true residual multiplied by the damping weight: damping then shrinks the reported error while moving the iteration *away* from the fixed point. Removing the damping and measuring correctly, the baseline reaches 5×10^{-5} on the value function and 8×10^{-9} on the price, and the two-bond economies of Section 6 meet the same joint criterion within a few hundred iterations. Removing the damping was what made the second of these possible: with it in place those solves had run to their iteration cap without converging.

The commitment economies need one refinement of the algorithm, not of the model. In those economies the bond price and the investment rule are jointly determined: the smoothed choice probabilities depend on the price schedule and the price schedule depends on those probabilities, so successive approximation makes the two chase each other. The value function still converges to machine precision, $\|T(V) - V\| \approx 2 \times 10^{-15}$, while the price residual stalls near 4×10^{-2}; the payoff is insensitive to the residual movement, but the price itself never settles.

Holding the policy fixed removes the coupling. Given the choice probabilities and the default decision, the pricing equation is linear in q,

$$q \;=\; R^{-1}\,\mathbb{E}\Big[(1-d')\big(\kappa + (1-\delta)\,\mathbb{E}\,q''\big)\Big],$$

a contraction with modulus $(1-\delta)/R = 0.878$ at our parameters, so it has a unique fixed point. We therefore solve the model as usual and then iterate the pricing equation alone, with the converged policy held fixed, until the price residual meets the same 10^{-5} criterion used everywhere else in the paper. In practice the polished residual ranges from 1.6×10^{-13} to 9.8×10^{-6} across the economies we solve, always inside that criterion.

The resulting price is an equilibrium object rather than an approximation to one, because

it must imply the policy from which it was computed. We verify this by restarting the full solver from the polished price: it moves q by 2×10^{-14} and leaves the value residual at 2×10^{-15}, so the policy implied by the polished price is the policy that generated it. Every economy in the paper is therefore solved at the same baseline taste-shock scale $\rho_b = 5 \times 10^{-3}$ and meets the same convergence criterion in both the value and the price.

Two details of the discretization are worth stating because the portfolio margin is close to indifferent. Among green shares whose value is within 10^{-5} of the best at a given (a', b'), we select the smallest, so that a tie is resolved by a stated refinement rather than by an arbitrary choice. The band and the converged price residual are of the same order, between 8×10^{-6} and 1×10^{-5}, so the refinement is doing real work: it picks among shares the solver cannot separate, rather than overriding a difference the solver has resolved. Because it always picks the smallest of those shares, the green shares we report are a lower bound on the share the government is willing to hold. And the debt and adaptation choices are discretized jointly, so a single index carries the pair (a', b') and the composition choice is made against the same grid as the investment choice.

E Debt-Financed Adaptation: Full Derivation

In this appendix, we derive the full optimality condition for debt-financed adaptation discussed in Section 4.3. Fix a consumption level $\bar{c}$ and the locus of policies (b', a') defined in equation (28). Totally differentiating (28) along the locus gives the debt the government must issue per unit of adaptation, $db'/da' = [f'(i) - q_{a'}n]/[q_{b'}n + q]$ with $n = b' - (1-\delta)b$. Along the locus the government's objective is $\beta\,\mathbb{E}V(b', a', y')$, and optimality requires its total derivative to vanish. The derivative with respect to b' is the right-hand side of (19) with the sign reversed, and the derivative with respect to a' is the right-hand side of (20), which carries both the protection channel, through y', and the surviving stock, through the state a' directly. Write $m(a', \gamma') = -h_a(a', \gamma') = (\gamma'^2/\bar{\gamma}) \exp(-a'\gamma'/\bar{\gamma})$ for the marginal protection of the technology (2), $\mathcal{V}_y^R = \partial_y V^R(b', a', y')$ for the marginal value of income under repayment, and $\Delta_y = -\partial_y \hat{\nu}(b', a', y')$ for its excess over the marginal value of income in default, so that $\mathcal{V}_y = \mathcal{V}_y^R - d(\chi')\,\Delta_y$. The condition is

$$\mathbb{E}\left[(1-d(\chi'))u'\left(c^{R'}\right)(\kappa+(1-\delta)q')\right]\left[\frac{f'(i)-q_{a'}(b'-(1-\delta)b)}{q_{b'}(b'-(1-\delta)b)+q}\right] =$$

$$\pi\left(\mathbb{E}[m]\,\mathbb{E}\left[\mathcal{V}_y^R\right]+\mathbb{COV}\left(\mathcal{V}_y^R, m\right)\right)$$

$$-\pi\left(\mathbb{E}[\Delta_y]\left(\mathbb{E}[m]\,\mathbb{E}[d(\chi')]+\mathbb{COV}(m,d(\chi'))\right)+\mathbb{COV}\left(\Delta_y,\, m\,d(\chi')\right)\right)+(1-\delta_a)\,\mathbb{E}\left[u'(c')f'(i')\right] \tag{33}$$

where every expectation on the right-hand side is conditional on a disaster next period. Under the linear technology $h(a,\gamma)=\gamma(1-a)$ the marginal protection is $m=\gamma'$ and the covariances are with the intensity itself.

The left-hand side of (33) represents the marginal utility cost of debt-financed adaptation. The first factor is the expected marginal cost of servicing one additional unit of debt, weighted by the probability of repayment and the bond payoff $\kappa+(1-\delta)q'$. The term in brackets is db'/da', the additional debt the government must issue to finance one unit of adaptation while holding consumption constant. This ratio depends on how prices respond to both debt and adaptation: if bond prices rise with adaptation ($q_{a'}>0$), the net cost of adaptation falls and less additional borrowing is needed.

The right-hand side decomposes the benefits. The first line captures the direct income effect: adaptation reduces disaster losses, and its value depends on the expected marginal protection and on the covariance between the marginal value of income and the severity of the disaster. When disasters are negatively correlated with the business cycle, the marginal value of the income they destroy is higher, which increases the incentive to invest in adaptation. The second line captures the insurance component related to default: in default the marginal value of income is lower by Δ_y, so if disasters are positively correlated with default episodes, the government has weaker incentives to invest in adaptation because default provides partial consumption insurance; the last covariance accounts for how adaptation changes that gap across states. The final term is the value of the stock that survives into the next period.

E.1 Proof of Lemma 1 and the Threshold in the Full Model

Proof of Lemma 1. Under the assumptions of the lemma the default decision has no cost shock, $\nu \equiv 0$, so $V(\chi) = \max\{V^R(b,a,y), V^D(a,y)\}$ and the government defaults when $\hat{\nu}(b,a,y) = V^D(a,y) - V^R(b,a,y)$ is positive; a defaulting government is excluded for good, $\theta = 0$; and from tomorrow on income equals the constant $\bar{y}$, so no disaster occurs and adaptation has no value: with $p_a > 0$ neither regime invests, both carry the stock $(1-\delta_a)a$ forward, and that stock affects nothing. Consumption today is $c^R = y - \kappa b + q(b')\,n$ under repayment and $c^D = y - D(y)$ under default. Write $W^R(b')$ for the continuation value of a government that repays today and carries debt b' into a future with constant income, and W^D for the continuation value of a government excluded for good; neither depends on today's income or on the stock, and neither does the price $q(b')$, which lenders set from tomorrow's deterministic default decision. Because a repaying government may still default tomorrow, $W^R(b') \geq W^D$ for every b'.

Step 1: consumption under repayment does not exceed consumption under default. At a state where the government defaults, $V^D(a,y) = u(c^D) + \beta W^D \geq V^R(b,a,y) = u(c^R) + \beta W^R(b'_R) \geq u(c^R) + \beta W^D$, so $c^D \geq c^R$. This is the argument of Arellano (2008): a plan that raised consumption above its level in default, together with the option to default later, would dominate default.

Step 2: the threshold falls with income. Since W^R, W^D, and q do not depend on y, income enters only through c^R and c^D, and the envelope theorem gives $\partial_y V^R = u'(c^R)$ and $\partial_y V^D = u'(c^D)\,[1 - D'(y)]$. Under the cost (9), $D'(y) = 0$ where the cost is inactive and $D'(y) = \eta_1 + 2\eta_2 y > \eta_2 y > 0$ where it is active, so by Step 1

$$\partial_y \hat{\nu}(b,a,y) \;=\; u'(c^D)\left[1 - D'(y)\right] - u'(c^R) \;\leq\; u'(c^D) - u'(c^R) \;\leq\; 0$$

at every state where the government defaults. This is the concavity step in Arellano (2008): an extra unit of income is worth more in the regime with lower consumption.

Step 3: a larger stock leaves more income. Income in the disaster state is $y(a) = \rho\, y_{t-1} - \xi\, h(a,\gamma) + \epsilon$ with $\xi > 0$, so $dy/da = \xi\,(-h_a(a,\gamma)) > 0$; under the technology (2) this is $\xi\,(\gamma^2/\bar{\gamma})\exp(-a\gamma/\bar{\gamma})$.

Conclusion. The stock in place enters the default decision only through the income it leaves, so by the chain rule, at every state where the government defaults,

$$\frac{d}{da}\,\hat{\nu}\big(b,a,y(a)\big) \;=\; \xi\,\big(-h_a(a,\gamma)\big)\,\partial_y \hat{\nu}\big(b,a,y(a)\big) \;\leq\; 0.$$

Suppose the government defaults at a_H, so that $\hat{\nu}(b,a_H,y(a_H)) > 0$, and let a^* be the

infimum of the stocks $a \le a_H$ such that $\hat{\nu}(b, \cdot, y(\cdot)) > 0$ on $[a, a_H]$. On $[a^*, a_H]$ the derivative is nonpositive, so $\hat{\nu}(b, a, y(a)) \ge \hat{\nu}(b, a_H, y(a_H)) > 0$ there, and by continuity the same holds on a neighborhood below a^* unless a^* is the lowest admissible stock. Hence the government defaults at every $a_L < a_H$. □

Sufficient conditions for the threshold to fall with income. In the full model, with recurring disasters, persistent income, re-entry, and a stochastic default cost, the continuations depend on y and Step 2 of the proof acquires a continuation term. Write $m(\hat{\nu}) = \mathbb{E}_\nu \min\{\hat{\nu}, \nu\}$, so that the expected value of the default option is $\mathbb{E}_\nu V(b, a, y, \nu) = V^D(a, y) - m\big(\hat{\nu}(b, a, y)\big)$, with $m'(\hat{\nu}) = 1 - F(\hat{\nu})$ the probability of repayment. Substituting into the two Bellman equations, a defaulting government re-enters with no debt, so its continuation is $V^D(a', y') - \theta\, m\big(\hat{\nu}(0, a', y')\big)$, and

$$\hat{\nu}(b, a, y) \;=\; u(c^D) - u(c^R) \;+\; \beta\, \mathbb{E}_y\big[m\big(\hat{\nu}(b'_R, a'_R, y')\big)\big] \;-\; \beta\theta\, \mathbb{E}_y\big[m\big(\hat{\nu}(0, a'_D, y')\big)\big] \\ +\; \beta\, \mathbb{E}_y\big[V^D(a'_D, y') - V^D(a'_R, y')\big],$$

evaluated at the optimal choices (b'_R, a'_R) under repayment and a'_D under default. The last term is the one the durable stock adds: the two governments carry different stocks into autarky. Under the level representation $y' = \rho y - \xi' h(a', \gamma') + \epsilon'$ the income process is a location shift, so $\partial_y \mathbb{E}_y[W(y')] = \rho\, \mathbb{E}_y[W'(y')]$ for any differentiable W. Differentiating with the envelope theorem, and using that a government with no debt repays, so $m'\big(\hat{\nu}(0, \cdot)\big) = 1$,

$$-\;\partial_y \hat{\nu}(b, a, y) \;=\; g(y) \;+\; \beta\rho \Big\{ \mathbb{E}_y\big[\Pr\{\text{repay}'\}\,\big(-\partial_{y'}\hat{\nu}(b'_R, a'_R, y')\big)\big] \;-\; \theta\, \mathbb{E}_y\big[-\partial_{y'}\hat{\nu}(0, a'_D, y')\big] \\ -\; \mathbb{E}_y\big[\partial_{y'} V^D(a'_D, y') - \partial_{y'} V^D(a'_R, y')\big] \Big\}, \qquad g(y) \;\equiv\; u'(c^R)\big[1 + q_y\, n\big] - u'(c^D)\big[1 - D'(y)\big],$$

where $n = b' - (1 - \delta)b$ is net issuance, q_y the derivative of the price schedule with respect to income, and $D(y) = \max\{0,\ \eta_1 y + \eta_2 y^2\}$ the output cost of default. Three conditions make the right-hand side positive at every state:

(i) $g(y) > 0$: a unit of income is worth more in repayment, where it raises consumption one for one and is sold into a better price, than in default, where the cost D absorbs part of it. With constant relative risk aversion this reads $(c^D/c^R)^\sigma > [1 - D'(y)]/[1 + q_y n]$. Under a cost that caps income in default, as in Arellano (2008), $u'(c^D)[1 - D'(y)] = 0$ above the cap and (i) is immediate.

(ii) $\frac{\sigma}{c^R}\,[1+q_y\,n]\,[\kappa+(1-\delta)q] \;\geq\; (1-\delta)\,q_y$: the period objective is supermodular in debt and income, because $\partial_b c^R = -[\kappa+(1-\delta)q]$ and $\partial_y c^R = 1+q_y n$ give $\partial^2_{by}\,u(c^R) = |u''(c^R)|\,[\kappa+(1-\delta)q]\,[1+q_y n] - u'(c^R)\,(1-\delta)\,q_y$. Since neither continuation depends on b, this makes $-\hat{\nu}$ supermodular in (b,y) whenever its continuation is, and supermodularity gives $-\partial_y\hat{\nu}(b',a,y) \geq -\partial_y\hat{\nu}(0,a,y)$ for every $b' \geq 0$.

(iii) $\Pr\{\text{repay}\}\cdot\inf_{a,y}\big(-\partial_y\hat{\nu}(0,a,y)\big) \;\geq\; \theta\,\sup_{a,y}\big(-\partial_y\hat{\nu}(0,a,y)\big) + \sup_y \operatorname{osc}_a\,\partial_y V^D(a,y)$, where osc_a is the variation across the stock: the government must be likelier to repay than to regain market access, up to the dispersion of the marginal value of income at zero debt and an allowance for the stock. The allowance is small because the direct dependence of $\partial_y V^D$ on the stock runs through $\partial_y c^D$ and is scaled by the resource value of the surviving stock, $(1-\delta_a)f'(i)$.

Under (i)–(iii), an induction on the Bellman operator delivers the property. Let $\mathcal{S}$ be the set of bounded continuous value functions whose threshold $\hat{\nu}$ is nonincreasing in y and whose $-\hat{\nu}$ is supermodular in (b,y) at every stock; $\mathcal{S}$ is closed under the supremum norm. For a candidate in $\mathcal{S}$, (ii) and (iii) make the bracket above nonnegative, so the updated $-\partial_y\hat{\nu}$ is at least $g>0$, and (ii) makes the updated $-\hat{\nu}$ supermodular; the operator maps $\mathcal{S}$ into itself. It is a contraction of modulus β, so its unique fixed point lies in $\mathcal{S}$ and satisfies $-\partial_y\hat{\nu} \geq g > 0$. The conditions are on equilibrium objects. On the calibrated solution the threshold falls with income at every state where default is a live option.

E.2 Proof of Lemma 2

Proof. The standard bond's price satisfies the recursion

$$q_s(\chi) \;=\; \frac{1}{R}\,\mathbb{E}\big[\,(1-d(\chi'))\,\big(\kappa+(1-\delta)\,q_s(\chi')\big)\;\big|\;\chi\big]\,.$$

The green bond pays the same flows plus the step-up, so its price satisfies

$$q_g(\chi) \;=\; \frac{1}{R}\,\mathbb{E}\big[\,(1-d(\chi'))\,\big(\kappa+\phi\,v(\chi')+(1-\delta)\,q_g(\chi')\big)\;\big|\;\chi\big]\,.$$

We wish to characterize the difference. Subtracting the first recursion from the second,

$$q_g(\chi)-q_s(\chi) \;=\; \frac{1}{R}\,\mathbb{E}\big[\,(1-d(\chi'))\,\big(\phi\,v(\chi')+(1-\delta)\,\big(q_g(\chi')-q_s(\chi')\big)\big)\;\big|\;\chi\big]\,.$$

The right-hand side is an affine operator in the difference q_g-q_s whose linear part has modulus $(1-\delta)/R<1$, so it is a contraction on the space of bounded functions and the

difference is its unique bounded fixed point. Iterating it forward T times leaves a remainder bounded by $\big((1-\delta)/R\big)^T \|q_g - q_s\|_\infty$, which vanishes as $T \to \infty$; the difference therefore equals the infinite discounted sum displayed in the lemma. The sum is well defined and the expectation and summation may be exchanged by monotone convergence, since $\phi \geq 0$, $v \in \{0,1\}$ and $d \in \{0,1\}$ make every term nonnegative. Nonnegativity of the terms gives $q_g(\chi) \geq q_s(\chi)$ directly, and the sum vanishes if and only if each term does, that is if and only if $\phi\, v(\chi_t) = 0$ almost surely at every date at which $\prod_{j \leq t}(1 - d(\chi_j)) = 1$. For $\phi > 0$ this is the condition stated. □

E.3 The Composition Margin

Fix a state $\chi = \{b, a, y, \nu, \vartheta\}$ in repayment and the choices b' and a'. Consumption is

$$c \;=\; y - \kappa b - \phi\, v(\chi)\, \vartheta b - f(i) + q_g(\vartheta')\big[\vartheta' b' - (1-\delta)\vartheta b\big] \\ + q_s(\vartheta')\big[(1-\vartheta')b' - (1-\delta)(1-\vartheta)b\big],$$

where $v(\chi)$ indicates that the current investment falls below the rule, and the prices depend on the composition ϑ' through next period's default, violation and continuation. Differentiating with respect to ϑ',

$$\frac{\partial c}{\partial \vartheta'} \;=\; (q_g - q_s)\, b' + n_g\, \partial_{\vartheta'} q_g + n_s\, \partial_{\vartheta'} q_s \;=\; P\, b' + n\, \partial_{\vartheta'} q_s + n_g\, \partial_{\vartheta'} P,$$

using $q_g = q_s + P$. Next period's government inherits $\vartheta = \vartheta'$. The share enters $V^R(b, a, y, \vartheta)$ only through the budget, so by the envelope theorem, at that government's own optimal choices,

$$\begin{aligned} \partial_\vartheta V^R(b, a, y, \vartheta) \;&=\; -\,u'(c^R)\, b\, \big[\phi\, v(\chi) + (1-\delta)\big(q_g - q_s\big)\big] \\ &=\; -\,u'(c^R)\, b\, \big[\phi\, v(\chi) + (1-\delta)P(\chi)\big], \end{aligned}$$

while V^D does not depend on ϑ because default wipes both bonds. Hence

$$\partial_{\vartheta'}\, \mathbb{E}\big[V(\chi')\big] \;=\; -\,b'\, \mathbb{E}\Big[\big(1 - d(\chi')\big)\, u'(c^{R'})\, \big(\phi\, v(\chi') + (1-\delta)P(\chi')\big)\Big],$$

the term at the boundary of the default set vanishing because the values of repayment and default are equal there. Setting $u'(c)\, \partial c/\partial \vartheta' + \beta\, \partial_{\vartheta'} \mathbb{E}[V(\chi')] = 0$ gives (30). The recursion

in the proof of Lemma 2, $P(\chi) = R^{-1}\,\mathbb{E}[(1 - d(\chi'))(\phi\, v(\chi') + (1-\delta)P(\chi'))]$, combines the premium with the right-hand side into the valuation wedge stated in the text. The same envelope argument shows that V^R is weakly decreasing in ϑ at every state, which is the signed channel of Section 6: a greener legacy weakly raises next period's default probability, whatever it does to investment.

F The Commitment Economy and the Adaptation Rule

The covenant of Section 6 verifies investment against a rule $\bar{a}'(y, b, a)$: the investment the government would choose if it could commit to the investment margin. This appendix defines that economy and how the rule is computed.

The economy commits the government to the investment margin alone and leaves borrowing discretionary. Its consumption equivalent, 0.050 percent, is the value of commitment that the paper measures the green bond against: the value of closing the investment wedge alone. It is not an upper bound on what a contract written on adaptation can deliver, for the reason given at the end of this appendix.

The rule economy (commitment to investment alone). This is the baseline model with one change to the budget in repayment. Following the marking approach of Hatchondo et al. (2016), legacy debt is valued at a reference price:

$$c + f(i) \;=\; y - \kappa b + q(y, b', a')\, b' - Q_A(y, b, b')\,(1 - \delta)\, b,$$

with y post-disaster income as in the main text, where

$$Q_A(y, b, b') \;=\; \sum_{a'} \Pr^{\text{inc}}\big(a' \mid y, b, b'\big)\; q(y, b', a')$$

is the expected issue price at the current state conditional on the chosen debt b', the expectation being over the investment the incumbent policy attaches to that debt choice. Because the surviving stock is marked to Q_A, a government that deviates on investment at a given b' collects the revaluation of the entire stock, new and legacy: the investment wedge $H_{a'}$ of Section 4.2 closes. The borrowing margin is not left exactly Markov. Its first-order condition differs from (19) by $(1 - \delta)\, b\, [\partial_{b'} Q_A - q_{b'}]$, the response of the mark to debt through the incumbent's investment response, so the economy closes the investment wedge at the incumbent's investment response to debt and holds the dilution wedge on borrowing fixed only up to that term. We use it to define the rule and as the benchmark for closing $H_{a'}$, not as an exact decomposition of the two wedges. On the equilibrium path Q_A equals the realized issue price, so the contract transfers no resources; only the incentives change.

The rule. The rule $\bar{a}'(y, b, a)$ is the equilibrium investment policy of this economy, computed on the same grids as the baseline, so the covenant can compare the government's choice to the rule state by state. The rule asks for more investment where commitment would deliver

more, at high debt with active disasters (panel b of Figure 6). In the maturity experiment of Section 7, the rule economy is re-solved at each δ, so the rule is specific to the maturity. This economy is worth 0.050 percent of consumption (median 0.052): it changes the mean adaptation stock by 9.0 percent and moves the mean spread from 2.02 to 1.92 percent, but leaves the borrowing margin Markov. It is therefore not an upper bound on what a covenant can achieve, and the green bond of Section 6 exceeds it at every maturity we solve. The reason is the margin this economy holds fixed: marking the legacy stock to Q_A changes the investment incentive and nothing else, whereas a coupon that steps up on violation is priced into the debt itself and so acts on borrowing as well. The rule economy measures the value of closing $H_{a'}$ alone; the covenant closes part of $H_{a'}$ and part of $H_{b'}$ at the same time.

The value of commitment. We measure it with the consumption equivalent of Section 5. Writing V^M for the value in the Markov equilibrium and V^C for the value under the contract, both evaluated at the same state, the equivalent λ solves $(1+\lambda)^{1-\sigma}V^M = V^C$, which at $\sigma = 2$ is

$$\lambda \;=\; \frac{V^M}{V^C} \;-\; 1,$$

positive when the contract is better. Averaged over the Markov ergodic distribution, the rule economy is worth 0.050 percent of consumption (median 0.052). The same statistic, re-solved at each maturity, is what Section 7 reports as the value of commitment. The solution uses the algorithm of Appendix D, with Q_A updated jointly with the price in each iteration and the price polish described there.