

Focusing attention, defocusing policy: climate shocks and the attention– policy translation failure

Giorgos Galanis, Giorgio Ricchiuti and Ben Tippet

July 2026

**Grantham Research Institute on
Climate Change and the Environment
Working Paper No. 444**

ISSN 2515-5717 (Online)

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training. The Institute is funded by the Grantham Foundation for the Protection of the Environment and a number of other sources. It has five broad research areas:

1. Climate change impacts and resilience
2. Cutting emissions
3. Financing a better future
4. Global action
5. Protecting the environment

More information about the Grantham Research Institute is available at: www.lse.ac.uk/GranthamInstitute

Suggested citation:

Galanis G, Ricchiuti G and Tippet B (2026) *Focusing attention, defocusing policy: climate shocks and the attention-policy translation failure*. Grantham Research Institute on Climate Change and the Environment Working Paper 444. London: London School of Economics and Political Science.

DOI: 10.21953/researchonline.lse.ac.uk.00140244

This working paper is intended to stimulate discussion within the research community and among users of research, and its content may have been submitted for publication in academic journals. It has been reviewed by at least one internal referee before publication. The views expressed in this paper represent those of the author[s] and do not necessarily represent those of the host institutions or funders.

Focusing Attention, Defocusing Policy: Climate Shocks and the Attention–Policy Translation Failure*

Giorgos Galanis ^{a b} Giorgio Ricchiuti ^c Ben Tippet ^d

^a *Queen Mary University of London*; ^b *London School of Economics and Political Science*; ^c *Università degli Studi di Firenze*; ^d *King's College London*

Abstract

This paper develops a diagnostic-expectations theory of climate focusing events. A disaster is diagnostic evidence about a latent hazard state, so voters overweight it when forming beliefs about future climate risk. This raises the perceived return to mitigation. But the same disaster destroys current output and raises reconstruction needs, tightening the resource constraint under which policy is produced. The paper's central claim is not that focusing events fail to focus attention; rather, attention may fail to translate into policy. In a political-economy model with diagnostic expectations and risk-sensitive (Hansen–Sargent) preferences, we decompose the disaster response of mitigation into a diagnostic focusing channel and a resource-defocusing channel. Mitigation falls when the latter dominates. The model also yields a non-monotone role for disaster severity: more severe shocks raise both the value of mitigation and the resource squeeze, but in the upper tail the resource squeeze can compound while the marginal value of further probability reduction saturates. Risk-sensitive preferences generate this saturation effect: when perceived disaster probability is already high, stronger diagnosticity need not always increase mitigation. Evidence from the 1990 Western European windstorm cluster illustrates the mechanism.

JEL Codes: D72, D91, Q54, Q58

*We thank Antonio Cabrales, Nicola Gennaioli, Iraklis Kollias, Yannis Papadakis, Elias Papaioannou, Gregor Singer, Boromeus Wanengkirtyo and Kostas Zachariadis and participants in WEHIA 2025, the RES Conference 2026 and seminars and workshops in London, Florence and Pisa. The usual disclaimer applies.

1 Introduction

Climate disasters will become more frequent and more intense over the coming decades (IPCC, 2022). Their economic costs are already large: USD 280 billion in 2023 alone, with annual average losses of USD 217 billion over 2014–2023 (Swiss Re Institute, 2024; Kahn, 2005), and these losses are amplified by network effects through global supply chains (Barrot and Sauvagnat, 2016). How disasters shape the political economy of mitigation policy is therefore a first-order question for the low-carbon transition. The usual intuition is that disasters should accelerate action because they are *focusing events*: they make future climate damages vivid, raise salience, and move voters toward climate policy (Birkland, 1998). Yet the movement from attention to legislation is not automatic. The central claim of this paper is that disasters can focus beliefs while defocusing policy output when post-disaster resource constraints bind.

The paper develops a diagnostic-expectations theory of this attention–policy translation failure. In Birkland’s framework, focusing events matter because they are sudden, harmful, and symbolically linked to a policy problem. We model this behavioral mechanism using diagnostic expectations (N. Gennaioli and Shleifer, 2010; Bordalo, N. Gennaioli, and Shleifer, 2018; Bordalo, N. Gennaioli, and Shleifer, 2022). A realised disaster is not treated as just another draw from a known distribution; it is overweighted as evidence about a latent hazard state. Formally, a disaster raises the perceived future hazard state, and therefore the perceived probability of future disaster losses, by more than under the rational conditional forecast. Diagnostic expectations are the paper’s formal microfoundation for focusing events.

The theoretical question is then: when does diagnostic focusing become mitigation policy? The answer depends on two opposing channels. The *diagnostic focusing channel* raises demand for mitigation. A disaster increases the perceived disaster probability; mitigation lowers that probability; and, under risk-sensitive preferences, voters value mitigation not only because it lowers expected losses but also because it narrows the future-consumption lottery. The *resource-defocusing channel* works in the opposite direction. The same disaster destroys current output and creates reconstruction needs, raising the shadow cost of the resources, administrative effort, and political attention required to produce mitigation policy. The model shows that these two channels can move in opposite directions: attention rises mechanically, but mitigation falls whenever the resource-defocusing channel dominates the diagnostic focusing channel.

The analysis delivers four theoretical results. First, a disaster is always a focusing event in beliefs: it raises the perceived hazard state and the perceived disaster probability, with diagnosticity scaling the belief response. Second, the optimal mitigation response can be decomposed into a diagnostic focusing term, a resource-defocusing term, and a reconstruction-side signal term. Third, the model characterises the defocusing region in which mitigation declines despite heightened perceived climate risk. This result also clarifies the role of disaster severity. More severe disasters raise the stakes of future climate losses, which can increase the value of mitigation, but they also destroy more current resources and intensify reconstruction needs. Severity therefore does not mechanically imply more mitigation. The upper-tail prediction is that especially severe shocks are prone to attention–policy translation failure when the resource squeeze rises faster than the marginal value of additional risk reduction. The role of diagnosticity is similarly sign-contingent: when the mitigation-side diagnostic bracket is positive, stronger diagnosticity makes defocusing less likely; when the saturation term dominates, stronger diagnosticity can reinforce defocusing. Fourth, in the dynamic extension, persistence is governed by the eigenvalues of the capital–hazard system. A short-lived belief shock can generate a long-lived policy shortfall when the resource/capital component is more persistent than the belief component.

Risk-sensitive preferences are used to discipline the mitigation-demand side of the mechanism. A disaster raises the perceived probability of the bad branch of the future-consumption lottery. Mitigation lowers that probability. Expected utility values this as expected loss avoided. Risk-sensitive preferences add an insurance motive: mitigation is valuable because it reduces the dispersion of continuation utility. This also creates a saturation effect. As perceived disaster probability rises, the marginal value of trimming it further can fall. The implication is that more salience does not necessarily imply more mitigation. Diagnostic expectations explain why disasters focus beliefs; risk-sensitive preferences explain why the belief shock has a rich, and sometimes non-monotone, effect on mitigation demand; resource constraints explain why policy can nevertheless move in the opposite direction.

To empirically motivate the model, we provide an illustration of a case in which an extreme weather shock was followed by an increase in public concern about global warming but was not translated into greater mitigation policy. In 1990, Western Europe experienced an unprecedented cluster of extreme weather events, with eight major storms striking the continent in quick succession. Finland, the Netherlands, Luxembourg, and Denmark were particularly affected, experiencing the largest increase in storm

frequency relative to their historical averages. Using an event-study design, we find that these four countries implemented fewer mitigation laws than the rest of Western Europe over the subsequent decade (1990–2000). While this analysis is neither a direct test nor a calibration of the theoretical model, it provides an empirical illustration of the attention–policy translation failure documented in the empirical literature reviewed in Section 2.

The contribution is primarily theoretical and connects three literatures. To the focusing-events literature (Birkland, 1998; Birkland, 2006), it gives a behavioral microfoundation: focusing events are diagnostic shocks to beliefs about a latent hazard state. To the diagnostic-expectations literature (N. Gennaioli and Shleifer, 2010; Bordalo, N. Gennaioli, and Shleifer, 2018; Bordalo, N. Gennaioli, and Shleifer, 2022; Afrouzi et al., 2023), it brings a new political-economy application in which overreaction to salient climate events need not generate more policy action. To the political economy of climate and disaster policy (Samuel Fankhauser, C. Gennaioli, and Collins, 2016; Sauquet, 2014; Carattini et al., 2023; de Silva and Tenreyro, 2021; Eskander and Sam Fankhauser, 2023; Galanis, Ricchiuti, and Tippet, 2025; Healy and Malhotra, 2009; Cohen and Werker, 2008; Gasper and Reeves, 2011; Garrett and Sobel, 2003; Besley and Persson, 2009; Besley and Persson, 2011), it identifies a resource-based mechanism through which disasters can crowd out long-run mitigation even as they heighten perceived climate risk.

The structure of the paper is as follows. Section 2 presents motivating facts from the 1990 Western European windstorm cluster. Section 3 develops the diagnostic focusing model, characterises the defocusing region, shows how disaster severity shapes that region, analyses persistence, and derives institutional-design implications. Section 4 concludes.

2 Motivating Facts

A useful way to motivate the mechanism is to distinguish shifts in *public demand* for climate action from changes in observable *mitigation policy* output. The literature points to a recurring pattern: extreme climate experiences can plausibly shift public demand, yet these shifts often do not translate into sustained mitigation policy production. This public demand–policy translation gap is the central motivation for the theoretical framework developed in the next section.

On the public demand side, a large literature suggests that salient climate disasters can move beliefs

and political preferences. For Europe, exposure to climate extremes is associated with higher environmental concern and greater Green voting, with the strength of these effects moderated by economic conditions (e.g., regional income), consistent with the view that the economic environment shapes whether climate experiences translate into expressed support for climate action (Hoffmann et al., 2022). Causal panel evidence reinforces this pattern across settings: U.S. disasters shift public perceptions of climate change (Sloggy et al., 2021), German flood experience strengthens climate-change beliefs (Osberghaus and Fugger, 2022), and flooding and heatwaves in England and Wales raise risk perception toward climate change (Lohmann and Kontoleon, 2023). Herrnstadt and Muehlegger (2014) push the link further toward observable political behaviour, showing that unusual local weather raises both climate-related Google searches and the propensity of U.S. congressional members to take pro-environment positions in roll-call votes. More broadly, Andre et al. (2024) document increased support for climate action in more vulnerable countries. Taken together, this evidence supports the premise that salient climate events can shift perceived climate threat and public demand—a force that corresponds to the diagnostic focusing channel in our theory.

A second set of findings, however, indicates that increased demand does not reliably translate into political supply. Using a large corpus of party communications across Europe, Wappenhans et al. (2024) show that, apart from short-lived responses by Green parties, mainstream parties do not increase attention to environmental issues following severe extreme weather events. This suggests a translation gap from public salience to agenda-setting and elite attention. More broadly, Otto et al. (2021) find that hazard frequency and severity are generally unassociated with improvements in disaster risk reduction policy across 85 countries once baseline policy levels and other covariates are accounted for. At the local level, when policy change occurs after extreme weather, it is typically oriented toward risk management and adaptation rather than mitigation; Giordono, Boudet, and Gard-Murray (2020) document scant evidence of mitigation-oriented policies proposed after events in a comparative analysis of 15 cases. Rowan (2022) also finds no statistically significant relationship between climate disasters and mitigation policy. Together, these findings are consistent with the possibility that post-disaster governance reallocates scarce policy effort toward urgent recovery and risk management tasks, crowding out mitigation even when perceived disaster probabilities rise.

2.1 An Empirical Illustration: The 1990 Western European Windstorm Cluster

To illustrate that extreme climate disasters do not necessarily lead to an increase in mitigation, we conduct an event study of the impact of a 1990 storm cluster in Western Europe on mitigation policies. Given the limitations of drawing definitive causal conclusions from a sample of only 16 Western European countries, these results should be interpreted simply as an illustration that motivates, rather than tests, the theoretical model.

The year 1990 was particularly severe for climate-related disasters in Western Europe, with eight major storms hitting the continent in quick succession and causing historically high levels of damage in several countries. Following common definitions of extremes in the climate-statistics literature (J. Hansen, Sato, and Ruedy, 2012; Rahmstorf and Coumou, 2011), we define an *extreme* year as one in which the annual count of climate-related disasters exceeds the country-specific mean by four standard deviations.¹ Using climate-related disaster data from the Emergency Events Database (EM-DAT) over the period 1980–2020, we find that this threshold is exceeded only in 1990, and only in four countries: Finland, the Netherlands, Luxembourg, and Denmark. Among the remaining Western European countries (Austria, Belgium, France, Germany, Ireland, Italy, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom), the threshold is never exceeded,² making the 1990 episode highly unusual.

Exploiting the fact that the 1990 shock was unique to a small subset of Western European countries, we conduct a simple event study comparing the treated countries (Finland, the Netherlands, Luxembourg, and Denmark) with the remaining Western European countries³ over the subsequent decade (1990–2000). Data on domestic mitigation legislation are taken from the Climate Change Laws of the World database. A detailed description of the data, empirical methodology, and robustness tests is

¹Annual counts of climate-related disasters, rather than damages as a share of GDP, are used to identify the exogenous shock because damages from natural disasters are influenced by endogenous factors, such as historical adaptation and risk-management investments. By contrast, an unusually large number of extreme events occurring within a country in a given year is highly exogenous and unpredictable.

²The only exception is Spain in 2019, which we exclude because it falls at the end of the sample period and its subsequent effects cannot be analysed.


³Austria, Belgium, France, Germany, Ireland, Italy, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom.

provided in the Appendix. Here, we present only the main event-study result.

Figure 1 shows the evolution of the cumulative number of mitigation laws in the treated and control groups. Pre-treatment trends are identical, as neither group implemented mitigation laws between 1980 and 1989. From 1990 onwards, however, the treated countries exhibit a relative decline in the cumulative number of mitigation laws compared with the control group. This difference becomes statistically significant at the 10% level three years after the 1990 shock and remains significant over the subsequent decade.

These findings provide an empirical illustration, consistent with the literature discussed above, that extreme climate disaster shocks do not necessarily lead to an increase in mitigation policies and may, in this case, reduce them. The Appendix presents a range of robustness checks that support this result. Although it is not possible to estimate a comparable event-study specification using public concern about global warming because of data limitations, Figure A.7 in the Appendix shows that countries experiencing larger GDP damages from the 1990 storm cluster also tended to exhibit larger increases in concern about global warming between 1989 and 1991, using Eurobarometer data.

Figure 1: Impact of climate disaster shock on cumulative no. mitigation laws



Note: 1990 is identified as the only year where the number of climate damages 4σ above the country mean. Treated group = Finland, Netherlands, Luxembourg and Denmark. Control group = Austria, Belgium, France, Germany, Ireland, Italy, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom. 90% confidence intervals

3 Model

We now turn to the model. Diagnostic expectations enter from the start because they are the formal counterpart of the focusing-events mechanism. A disaster is a salient realisation that voters overweight when forecasting a latent hazard state. The model then asks whether this diagnostic belief movement translates into mitigation policy once the government faces a post-disaster resource constraint.

The model has two parts. A two-period framework develops the core mechanism and identifies the parameter region in which the resource-defocusing channel dominates the diagnostic focusing channel. An infinite-horizon version then treats persistence as a local propagation problem: once the shock has reduced resources, the mitigation shortfall follows the eigenmodes of the capital–hazard system and can outlast the belief shock.

3.1 Two-Period Model

Environment

An incumbent government chooses policy in period $t = 0$ to maximize its probability of reelection. Elections are decided by probabilistic voting (Lindbeck and Weibull, 1987; Dixit and Londregan, 1996; Besley and Persson, 2011): a representative swing voter supports the incumbent if the voter’s perceived utility under the incumbent weakly exceeds that under the challenger, up to an idiosyncratic ideology shock. The incumbent’s reelection probability is increasing in the swing voter’s perceived expected utility from the incumbent’s policy platform.

In period 0, a climate disaster may occur: $D_0 \in \{0, 1\}$, with a probability determined by the underlying hazard state introduced below. Having observed D_0 , the incumbent chooses mitigation $m_0 \geq 0$ and reconstruction $r_0 \geq 0$. Period 1 involves no decisions, only stochastic consumption based on the realized hazard state.

Production, Damage, and Capital Dynamics

Output in each period depends on capital K_t , with initial stock $K_0 > 0$ and productivity $A > 0$, and on disaster damages:

$$Y_t = AK_t(1 - \delta D_t), \quad \delta \in (0, 1). \quad (1)$$

The damage fraction δ captures disaster severity and is directly interpretable as the share of output destroyed by a disaster, matching the damages-to-GDP measure used in Section 2. This parameter is conceptually distinct from both the realised event D_t and the disaster probability $p(H_t)$. The event D_t says whether a disaster occurs, $p(H_t)$ says how likely disaster occurrence is, and δ says how costly the event is conditional on occurrence. Thus, severity does not shift beliefs directly in the baseline model; it changes the resource loss today and the utility loss associated with future disaster states. Capital evolves as

$$K_1 = (1 - d)K_0 + r_0, \quad (2)$$

with $d \in (0, 1)$ a depreciation rate. The dynamic extension allows mitigation to have a productive co-benefit through investment in clean capital.

Disaster Probability, the Hazard State, and Diagnostic Beliefs

A disaster in period t occurs with probability

$$\Pr(D_t = 1 \mid H_t) = p(H_t), \quad (3)$$

with $p : \mathcal{H} \rightarrow (0, 1)$ increasing. We maintain throughout that p is twice continuously differentiable and weakly convex on the relevant hazard-state domain: $p'(H) > 0$ and $p''(H) \geq 0$ wherever evaluated. Convexity is the substantive half of this regularity: it signs the mean piece of the diagnostic focusing channel below, and the concavity argument in Appendix B uses $p'' \geq 0$ directly. The variable H_t is a latent *hazard state*: a slow-moving index of the physical conditions that make extreme weather more likely, summarising fundamentals such as cumulative emissions, temperature anomalies and the associated hazard conditions. The disaster probability is the transformed object $p_t \equiv p(H_t)$. Voters

care about the probability distribution over future consumption induced by p_t , while policy acts on the underlying hazard state H_t . We therefore use the following terminology throughout:

Object	Notation	Meaning
Realised disaster	D_t	Binary disaster event.
Conditional severity	δ	Output loss conditional on $D_t = 1$.
Hazard state	H_t	Latent physical climate/hazard index.
Disaster probability	$p_t = p(H_t)$	Objective probability of $D_t = 1$.
Perceived hazard state	$\hat{H}_t$	Diagnostic belief about the hazard state.
Perceived disaster probability	$\hat{p}_t = p(\hat{H}_t)$	Probability entering welfare beliefs.
Mitigation	m_t	Policy effort that lowers the future hazard state.
Reconstruction	r_t	Policy effort that rebuilds productive capital.

In informal discussion we sometimes use “climate risk” as an umbrella term, but the formal distinction is between the hazard state H_t , the disaster probability p_t , the perceived disaster probability $\hat{p}_t$, and the conditional severity parameter δ . This distinction matters for the mechanism. The focusing-event channel operates through the realised event D_0 and the belief response it induces about $\hat{H}_1$ and $\hat{p}_1$. Severity governs the size of the loss conditional on such an event. A severe disaster can therefore make mitigation more valuable by raising the bad-state loss, while simultaneously making mitigation harder to finance by reducing current resources.

The hazard state evolves as

$$H_1 = \rho H_0 - \alpha m_0 + \psi_H D_0, \quad (4)$$

with $\rho \in (0, 1)$ persistence, $\alpha > 0$ the effectiveness of mitigation in improving the hazard state, and $\psi_H \geq 0$ the feedback from a realised disaster to the future hazard state, a reduced form for the serial clustering of hazard conditions. The dynamic extension in Section 3.5 adds an i.i.d. climate-variability shock.

Voters form beliefs about the future hazard state using diagnostic expectations (Bordalo, N. Gennaioli, and Shleifer, 2018; Bordalo, N. Gennaioli, and Shleifer, 2022) with diagnosticity parameter $\theta \geq 0$. Following the standard formulation, the perceived future hazard state is the rational conditional forecast

plus an overreaction term that weights the surprise component of news:

$$\hat{H}_1 \equiv \mathbb{E}_0^{DE}[H_1] = \rho H_0 - \alpha m_0 + \psi_H D_0 + \theta \psi_H (D_0 - \mathbb{E}[D_0]). \quad (5)$$

The first three terms on the right-hand side are the rational conditional expectation $\mathbb{E}_0^{RE}[H_1 | D_0]$; the final term is the diagnostic overweighting of the surprise $D_0 - \mathbb{E}[D_0]$. The perceived disaster probability, the object that enters the voter's welfare, is $\hat{p}_1 = p(\hat{H}_1)$. Setting $\theta = 0$ recovers the rational benchmark, in which a realised disaster still raises the perceived hazard state by ψ_H per unit through the conditional expectation. With $\theta > 0$, the same disaster raises the perceived hazard state by $(1 + \theta)\psi_H$ per unit when it exceeds its prior mean, and perceived disaster probability with it through $p(\cdot)$, generating excessive pessimism relative to RE. Where the argument is the perceived hazard state, we write $\kappa_1 \equiv p'(\hat{H}_1) > 0$ and $\kappa_2 \equiv p''(\hat{H}_1) \geq 0$ for the local slope and curvature of the probability function.

Proposition 1. *Holding policy fixed, a realised disaster raises the perceived future hazard state by*

$$\frac{\partial \hat{H}_1}{\partial D_0} = (1 + \theta)\psi_H,$$

and raises the perceived disaster probability by

$$\frac{\partial \hat{p}_1}{\partial D_0} = p'(\hat{H}_1)(1 + \theta)\psi_H > 0$$

whenever $\psi_H > 0$. For $\theta > 0$, the belief response is larger than the rational conditional response.

Because D_0 is binary, the derivatives in Proposition 1 are exact for $\hat{H}_1$, which is linear in D_0 , and should be read for $\hat{p}_1$ as the local response along the continuous embedding introduced before Theorem 2; the discrete effect has the same sign by monotonicity of p , since $\hat{p}_1(D_0 = 1) - \hat{p}_1(D_0 = 0) = p(\hat{H}_1(0) + (1 + \theta)\psi_H) - p(\hat{H}_1(0)) > 0$ whenever $\psi_H > 0$.

Proposition 1 is the paper's first wedge. Every disaster is a focusing event in beliefs, and diagnostic expectations strengthen that focusing. Defocusing is therefore not a failure of attention. It is a failure of translation from heightened perceived risk into policy output.

Resource Constraints

Consumption equals output net of investments:

$$C_0 = AK_0(1 - \delta D_0) - m_0 - r_0, \quad C_1 = AK_1(1 - \delta D_1), \quad (6)$$

where D_1 is stochastic with disaster probability $p(H_1)$ as in (3).

Preferences

The representative swing voter has risk-sensitive recursive preferences in the Kreps–Porteus tradition (Epstein and Zin, 1989; Weil, 1990). The incumbent places an effective political continuation weight $\beta_P \in (0, 1)$ on future utility, capturing electoral horizons, legislative myopia, or the probability that the incumbent internalises future policy benefits. To keep notation light, all formulas below write $\beta \equiv \beta_P$. Period felicity is aggregated through the exponential certainty equivalent,

$$V_t = u(C_t) + \beta \mathcal{R}_t[V_{t+1}], \quad \mathcal{R}_t[X] \equiv -\frac{1}{\Theta} \ln \mathbb{E}_t^{DE} [e^{-\Theta X}], \quad (7)$$

with $V_1 = u(C_1)$ at the terminal date, so that $V_0 = u(C_0) + \beta \mathcal{R}_0[u(C_1)]$. The parameter $\Theta \geq 0$ indexes risk sensitivity over continuation utility.⁴ As $\Theta \rightarrow 0$, $\mathcal{R}_t$ converges to the conditional expectation and (7) reduces to expected utility; $\Theta > 0$ corresponds to a preference for early resolution of uncertainty. This is the risk-sensitive preference family of Tallarini Jr (2000) and L. P. Hansen and Sargent (2008), and it coincides with Epstein–Zin utility at unit intertemporal elasticity of substitution. A second-order expansion of (7) in Θ delivers the mean–variance representation

$$V_0 \approx u(C_0) + \beta \mathbb{E}_0^{DE} [u(C_1)] - \frac{\beta \Theta}{2} \text{Var}_0^{DE} [u(C_1)],$$

familiar from log-linearisations of Epstein–Zin utility (L. P. Hansen, Heaton, and Li, 2008; Restoy and Weil, 2011), with Θ in the role of the wedge between risk aversion and the inverse IES. We work with

⁴ Θ has the dimension of inverse utility, so its magnitude depends on the normalisation of u . The identification $\Theta = \gamma_{EZ} - 1/\psi_{EZ}$ is exact in the unit-free log-linearised representation of Epstein–Zin utility; more generally Θ should be read as a reduced-form risk-sensitivity parameter rather than calibrated directly from estimates of (γ_{EZ}, ψ_{EZ}) .

the exact recursion (7) throughout and use the mean–variance form only for interpretation.

Period felicity is maintained to be CRRA with curvature $\sigma \geq 1$ throughout: $u(C) = C^{1-\sigma}/(1-\sigma)$ for $\sigma > 1$, and $u(C) = \ln C$ at $\sigma = 1$, understood as the CRRA limit under the standard normalisation $(C^{1-\sigma} - 1)/(1-\sigma)$; the additive constant is irrelevant under the translation-invariant aggregator (7). The restriction $\sigma \geq 1$ does two jobs below: it makes felicity unbounded below as consumption vanishes, which the existence argument in Appendix B uses to keep the optimum away from the resource boundary, and it signs the cross-state marginal-utility gap $\Delta u_K \geq 0$ defined below, the direction of the reconstruction-side channel. The relationship between the risk adjustment and the CRRA specification merits one remark. The aggregator $\mathcal{R}_t$ operates on continuation utility, while the period felicity $u(\cdot)$ remains a free choice. The two parameters are mutually independent: $\Theta = 0$ recovers expected utility with curvature σ unaffected, and the pairing of recursive risk adjustment with CRRA felicity is standard in the long-run-risks literature (Bansal and Yaron, 2004).

With binary D_1 and perceived probability $\hat{p}_1$, write $u_n \equiv u(AK_1)$ and $u_d \equiv u((1-\delta)AK_1)$ for felicity in the no-disaster and disaster states, $\Delta u \equiv u_d - u_n < 0$, and $|\Delta u| = u_n - u_d$ for the utility loss from a period-1 disaster. The certainty equivalent is then available in closed form:

$$\mathcal{R}_0[u(C_1)] = u(AK_1) - \Phi, \quad \Phi \equiv \frac{1}{\Theta} \ln \left[1 + \hat{p}_1 \left(e^{\Theta|\Delta u|} - 1 \right) \right], \quad (8)$$

where Φ is the entropic risk premium attached to disaster exposure. Two derived objects organise everything that follows. The *marginal value of probability reduction*,

$$\mathcal{M} \equiv \frac{\partial \Phi}{\partial \hat{p}_1} = \frac{e^{\Theta|\Delta u|} - 1}{\Theta [1 + \hat{p}_1 (e^{\Theta|\Delta u|} - 1)]} > 0, \quad (9)$$

is the welfare gain from trimming perceived disaster probability by one unit. It satisfies $\mathcal{M} \rightarrow |\Delta u|$ as $\Theta \rightarrow 0$ and, to first order in Θ , $\mathcal{M} = |\Delta u| [1 + \frac{\Theta}{2} (1 - 2\hat{p}_1) |\Delta u|]$: the expected loss avoided plus a variance-hedging premium. The *tilted probability*,

$$\pi \equiv \frac{\hat{p}_1 e^{\Theta|\Delta u|}}{1 + \hat{p}_1 (e^{\Theta|\Delta u|} - 1)} \in [\hat{p}_1, 1), \quad (10)$$

is the risk-adjusted weight on the disaster state when the voter values marginal resources: $\pi = \hat{p}_1$ under

expected utility, and π rises with Θ . Two properties drive the results below. First, $\mathcal{M} > 0$ everywhere: reducing perceived disaster probability always raises welfare, so the marginal benefit of mitigation is positive on the whole probability range. Second, $\partial\mathcal{M}/\partial\hat{p}_1 = -\Theta\mathcal{M}^2 \leq 0$, strictly for $\Theta > 0$ and all $\hat{p}_1 \in (0, 1)$: the marginal value of probability reduction declines as perceived disaster probability rises. This global saturation property is the exact counterpart of the variance-hedging intuition. In the mean–variance approximation the factor $(1 - 2\hat{p}_1)$ appears in the *level* of $\mathcal{M}$ —the hedging premium $\frac{\Theta}{2}(1 - 2\hat{p}_1)|\Delta u|^2$ is positive only for $\hat{p}_1 < 1/2$ —but not in its slope: the first-order approximation satisfies $\partial\mathcal{M}^{(1)}/\partial\hat{p}_1 = -\Theta|\Delta u|^2 < 0$ on the entire unit interval, matching the exact property.

Two further derived objects appear below: $\Pi \equiv \partial\pi/\partial\hat{p}_1 = \pi(1 - \pi)/[\hat{p}_1(1 - \hat{p}_1)] > 0$, the sensitivity of the tilted probability to perceived disaster probability, which tends to one as $\Theta \rightarrow 0$; and $\Delta u_K \equiv (1 - \delta)Au'((1 - \delta)AK_1) - Au'(AK_1) = Au'(AK_1)[(1 - \delta)^{1-\sigma} - 1] \geq 0$ for $\sigma \geq 1$, the gap in the marginal utility value of period-1 capital between the disaster and no-disaster states, which vanishes in the log limit.

Assumptions

Throughout the two-period analysis we maintain the regularity conditions introduced above: primitives lie in the domains stated at their introduction, felicity is CRRA with $\sigma \geq 1$, and the probability function is a twice continuously differentiable, increasing, and weakly convex map of the relevant hazard-state domain into $(0, 1)$. Two further assumptions discipline the policy problem; both hold at every exposure $s \in [0, 1]$ of the embedding (15) introduced below, whose endpoints $s = 0$ and $s = 1$ reproduce the no-disaster and disaster problems. For each s , write $V_0(\cdot; s)$ for the objective after the embedding and

$$\mathcal{F}(s) = \{(m_0, r_0) \in \mathbb{R}_+^2 : m_0 + r_0 < AK_0(1 - \delta s)\}$$

for the feasible set.

Assumption 1 (Strict concavity). At every $(m_0, r_0) \in \mathcal{F}(s)$ and every exposure $s \in [0, 1]$,

$$(\sigma C_0^{-\sigma-1} - \beta\Theta\alpha^2\kappa_1^2\mathcal{M}^2)(\sigma C_0^{-\sigma-1} + \mathcal{G}) > (\sigma C_0^{-\sigma-1} + \mathcal{W})^2, \quad (11)$$

where

$$\mathcal{G} \equiv \beta(1 - \pi)A^2 |u''(AK_1)| + \beta\pi(1 - \delta)^2 A^2 |u''((1 - \delta)AK_1)| + \beta\Theta\pi(1 - \pi)(\Delta u_K)^2 > 0,$$

$$\mathcal{W} \equiv \beta\alpha\kappa_1 \Pi \Delta u_K \geq 0$$

are the capital-side own curvature and the belief–capital cross-wedge.

Assumption 2 (Upper-contour interiority). For each exposure $s \in [0, 1]$, let

$$\mathcal{U}(s) = \{(m_0, r_0) \in \overline{\mathcal{F}(s)} : V_0(m_0, r_0; s) \geq V_0(0, 0; s)\}$$

be the relevant upper-contour set, understood after the compactness argument in the proof has excluded the boundary $C_0 = 0$. At every point of $\mathcal{U}(s)$ with $m_0 = 0$,

$$\beta\alpha\kappa_1 \mathcal{M} > u'(C_0), \tag{12}$$

and at every point of $\mathcal{U}(s)$ with $r_0 = 0$, $\beta\mathcal{B} > u'(C_0)$, where

$$\mathcal{B} \equiv (1 - \pi) Au'(AK_1) + \pi(1 - \delta) Au'((1 - \delta)AK_1)$$

is the tilted-mean marginal product of capital in utility terms, the marginal value of reconstruction.

Assumptions 1 and 2 are sufficient conditions for a unique interior optimum. Assumption 1 is the strict-concavity restriction that keeps the two-instrument problem well behaved: it rules out cases in which the saturation term in the probability-reduction motive makes the mitigation margin locally convex, and it bounds the belief–capital cross-wedge generated by reconstruction. Assumption 2 rules out corner solutions in which either mitigation or reconstruction is zero on the economically relevant upper-contour set. Assumption 2 is a standard interiority device; Assumption 1, by contrast, carries substantive content at an interior optimum. Combining (11) with the first-order conditions and taking

leading order ($\Theta \rightarrow 0$, $\delta \rightarrow 0$, $\Pi \rightarrow 1$) reduces it to the curvature bound

$$\frac{C_0}{K_1} < \frac{\sigma(2-\sigma)}{(\sigma-1)^2}, \quad (13)$$

derived in Appendix B: the right-hand side is non-positive for $\sigma \geq 2$, so Assumption 1 caps felicity curvature below two, and for $\sigma > 1 + \frac{\sqrt{2}}{2} \approx 1.71$ it requires a consumption–capital ratio below one. Risk sensitivity tightens rather than relaxes the restriction: the first factor in (11) must itself be strictly positive, and because (FOC-m) implies $\beta\Theta\alpha^2\kappa_1^2\mathcal{M}^2 = \Theta u'(C_0)^2/\beta$ exactly at an interior optimum, a necessary condition for (11) there is $\Theta < \sigma\beta C_0^{\sigma-1}$, an effective ceiling on risk sensitivity in the units of u (Appendix B). The economically relevant region for the mechanism is therefore $\sigma \in [1, 2)$ with moderate risk sensitivity; within it, Assumptions 1–2 allow the comparative statics to be stated using first-order conditions. Appendix B derives the exact Hessian terms, explains how Assumption 1 signs the substitution pattern used below, and provides these primitive interpretations.

The Incumbent’s Problem

Under probabilistic voting, maximising the probability of reelection is equivalent to maximising the swing voter’s perceived value (7). The incumbent therefore solves

$$\max_{m_0, r_0 \geq 0} V_0, \quad (14)$$

subject to (1)–(6), taking the diagnostic perception (5) as the relevant belief.

3.2 Results

Any interior solution of the incumbent’s problem satisfies two first-order conditions, derived in Appendix B:

$$u'(C_0) = \beta\alpha\kappa_1\mathcal{M}, \quad u'(C_0) = \beta\mathcal{B}.$$

The first equates the marginal consumption cost of mitigation to its marginal benefit: each unit of mitigation lowers the perceived disaster probability by $\alpha\kappa_1$, and each unit of probability reduction is

worth $\mathcal{M}$. The marginal value $\mathcal{M}$ embeds both the expected-utility motive and the precautionary one. At $\Theta = 0$ it is the expected utility loss avoided, $|\Delta u|$; for $\Theta > 0$ it adds a premium for the reduction in the dispersion of period-1 outcomes, the term $\frac{\Theta}{2}(1 - 2\hat{p}_1)|\Delta u|^2$ at first order. The second condition equates the marginal cost of reconstruction to the tilted-mean marginal product of capital $\mathcal{B}$ defined in Assumption 2: a risk-sensitive voter values next-period capital under the tilted probability $\pi \geq \hat{p}_1$, overweighting the disaster state, where marginal utility is higher when $\sigma > 1$. Because π moves with beliefs, the reconstruction margin responds to perceived disaster probability as well; this generates the reconstruction-side signal term in the decomposition below. Both marginal benefits are strictly positive on the whole probability range, with no restriction on the level of perceived disaster probability.

Theorem 1. *Under Assumptions 1 and 2, the incumbent's problem admits a unique interior solution (m_0^*, r_0^*) with $m_0^*, r_0^* > 0$. The solution is the global maximum of V_0 .*

We now turn to how mitigation responds to a disaster. Because D_0 is binary, a derivative with respect to D_0 is not defined; the object of interest is the discrete effect $m_0^*(D_0 = 1) - m_0^*(D_0 = 0)$. To characterise it, we embed the two realisations in a continuous family of problems indexed by an exposure parameter $s \in [0, 1]$, setting $D_0 = s$ in the period-0 constraint and in the belief equation:

$$C_0(s) = AK_0(1 - \delta s) - m_0 - r_0, \quad \hat{H}_1(s) = \rho H_0 - \alpha m_0 + \psi_H s + \theta \psi_H (s - \mathbb{E}[D_0]), \quad (15)$$

where $\mathbb{E}[D_0] = p(H_0)$ is the rational pre-realisation disaster probability, predetermined by the inherited state and hence constant along the path. The endpoints $s = 0$ and $s = 1$ reproduce the no-disaster and disaster problems exactly; interior values have a natural reading as partial exposure, with s the share of the economy hit by the shock. The maintained regularity conditions and Assumptions 1–2 hold at every $s \in [0, 1]$; in particular, the relevant hazard-state domain for the probability function is read uniformly along the path. Theorem 1 then applies at each s , and since $\Delta_H > 0$ the implicit function theorem makes the optimal policy path $s \mapsto (m_0^*(s), r_0^*(s))$ continuously differentiable, so that

$$\Delta m_0^* \equiv m_0^*(1) - m_0^*(0) = \int_0^1 \frac{dm_0^*(s)}{ds} ds. \quad (16)$$

Differentiating the first-order conditions along the embedding delivers the paper's central decomposition:

a diagnostic focusing channel, a resource-defocusing channel, and a reconstruction-side signal term.

Theorem 2. *Under the conditions of Theorem 1, maintained at every $s \in [0, 1]$, the marginal response of optimal mitigation along the embedding path is*

$$\frac{dm_0^*(s)}{ds} = \frac{1}{\Delta_H} [-V_{rr} \cdot S + V_{rr} \cdot O - V_{mr} \cdot O + V_{mr} \cdot S_r], \quad (17)$$

where $\Delta_H = V_{mm}V_{rr} - V_{mr}^2 > 0$ is the Hessian determinant and

$$S \equiv (1 + \theta)\psi_H \beta \alpha [\kappa_2 \mathcal{M} - \Theta \kappa_1^2 \mathcal{M}^2], \quad (18)$$

$$O \equiv -u''(C_0) \delta A K_0 > 0, \quad (19)$$

$$S_r \equiv \beta(1 + \theta)\psi_H \kappa_1 \Pi \Delta u_K \geq 0, \quad (20)$$

with $\Delta u_K \geq 0$ for $\sigma \geq 1$ and $\Pi > 0$ as defined in Section 3.1. All terms on the right-hand side are evaluated at $(m_0^*(s), r_0^*(s); s)$.

The same expression also signs the discrete disaster effect. If the bracketed term in (17) is strictly negative at every $s \in [0, 1]$, then the disaster effect satisfies $\Delta m_0^* < 0$ in (16); if it is strictly positive at every s , then $\Delta m_0^* > 0$.

The diagnostic focusing component S is the product of the belief shift $(1 + \theta)\psi_H$ and the response of mitigation's marginal benefit $\beta \alpha \kappa_1 \mathcal{M}$ to the perceived hazard state, and it has two pieces. The *mean piece*, $\kappa_2 \mathcal{M}$, is active when the probability function is convex: a higher perceived hazard state steepens $p(\cdot)$ locally, raising the probability reduction each unit of mitigation buys. The *saturation piece*, $-\Theta \kappa_1^2 \mathcal{M}^2$, is the decline in the marginal value of probability reduction as the disaster pushes perceived disaster probability up, the global property $\partial \mathcal{M} / \partial \hat{p}_1 = -\Theta \mathcal{M}^2$ of the certainty equivalent. Expanding $\mathcal{M}$ to first order in Θ recovers the mean-variance decomposition of S into a mean-channel piece $\beta \alpha \kappa_2 |\Delta u|$, a variance-hedging piece $\frac{\beta \Theta}{2} \alpha \kappa_2 (1 - 2\hat{p}_1)(\Delta u)^2$, positive for $\hat{p}_1 < 1/2$, and a negative variance-hedging piece $\beta \Theta \alpha \kappa_1^2 (\Delta u)^2$; the exact two-piece form supersedes that expansion. The resource-defocusing component O is strictly positive. The reconstruction-side diagnostic focusing component S_r satisfies $\Delta u_K \approx (\sigma - 1) \delta A u'(A K_1)$ at leading order: it is strictly positive for $\sigma > 1$ and vanishes identically in the log-felicity limit $\sigma = 1$, in which case $\Delta u = \ln(1 - \delta)$ is independent of K_1 and the

decomposition reduces to the pure two-channel form.

The interpretation of the four terms in (17) is straightforward. The first is the direct diagnostic focusing effect: since $V_{rr} < 0$, it raises mitigation when $S > 0$, reflecting the rise in perceived disaster probability transmitted through belief updating (the rational response, plus its diagnostic amplification). The second is the direct resource-defocusing effect: with $V_{rr} < 0$ and $O > 0$, it lowers mitigation as resources tighten. The third is a cross-effect: reconstruction adjusts in response to the output loss, with sign determined by V_{mr} . In our setting $V_{mr} < 0$, so this term is positive and partially offsets the direct resource-defocusing effect. The fourth is a reconstruction-side diagnostic focusing effect: the same rise in perceived disaster probability that fuels demand for mitigation also raises the tilted weight π on the disaster state, where marginal utility is higher when $\sigma > 1$, so the precautionary value of capital rises; resources are pulled toward reconstruction, and with $V_{mr} < 0$ and $S_r \geq 0$ the term $V_{mr} \cdot S_r$ weakly lowers mitigation. It vanishes under log felicity.

The internal structure of S is worth unpacking. Under linear probability ($\kappa_2 = 0$) the mean piece vanishes, leaving $S = -(1 + \theta)\psi_H\beta\Theta\alpha\kappa_1^2\mathcal{M}^2 < 0$. In this case the diagnostic focusing channel actually works *against* the precautionary motive: by pushing $\hat{p}_1$ up, a disaster lowers the marginal value of further probability reduction. Under expected utility ($\Theta = 0$), only the mean piece survives, $S = (1 + \theta)\psi_H\beta\alpha\kappa_2|\Delta u|$, which is positive when $\kappa_2 > 0$ and vanishes under linear p . In general the sign of S turns on the comparison of κ_2 with $\Theta\kappa_1^2\mathcal{M}$: writing $\Theta\mathcal{M} = (e^{\Theta|\Delta u|} - 1)/[1 + \hat{p}_1(e^{\Theta|\Delta u|} - 1)]$, the signal is negative precisely when this ratio exceeds κ_2/κ_1^2 . Note that even at $\theta = 0$ (RE) the diagnostic focusing channel does not vanish: a realised disaster still shifts beliefs via the conditional expectation, and the diagnostic parameter scales that baseline by a factor of $(1 + \theta)$.

Expected-utility benchmark. When $\Theta = 0$, the mitigation-side diagnostic focusing channel is

$$S^{EU} = (1 + \theta)\psi_H\beta\alpha\kappa_2|\Delta u|.$$

It is positive if $p''(\hat{H}_1) > 0$, zero if p is locally linear, and is amplified one-for-one by diagnosticity through the factor $(1 + \theta)$. Thus the basic attention-policy wedge does not rely on risk sensitivity: risk sensitivity enriches the diagnostic focusing channel by adding the saturation term (Appendix C).

The Defocusing Region

The decomposition delivers a sharp characterisation of when disasters reduce mitigation.

Proposition 2. *Under the conditions of Theorem 2, define the output-signal ratio*

$$\Lambda(s) \equiv \frac{O(V_{rr} - V_{mr})}{V_{rr}}, \quad (21)$$

and its adjusted form $\tilde{\Lambda}(s) \equiv \Lambda(s) + (V_{mr}/V_{rr}) S_r \geq \Lambda(s)$, with all objects evaluated along the embedding path at $(m_0^*(s), r_0^*(s); s)$. Then:

(i) At each $s \in [0, 1]$,

$$\frac{dm_0^*(s)}{ds} < 0 \quad \Longleftrightarrow \quad \tilde{\Lambda}(s) > S(s). \quad (22)$$

(ii) If $\tilde{\Lambda}(s) > S(s)$ at every $s \in [0, 1]$, then $\Delta m_0^* < 0$: the disaster lowers mitigation. Symmetrically, if $\tilde{\Lambda}(s) < S(s)$ at every s , then $\Delta m_0^* > 0$.

At the margin, defocusing occurs if and only if the resource-defocusing channel (scaled by reconstruction substitution), together with the reconstruction-side signal, exceeds the mitigation-side diagnostic focusing channel; by part (ii), the disaster effect itself is negative whenever this dominance holds along the whole path. The uniformity requirement has little bite. Between the endpoints, the resource-side coefficients move by order δ (through $C_0(s)$) and the belief-side coefficients by order $\kappa_1(1+\theta)\psi_H$ (through $\hat{p}_1(s)$), so if the inequality holds strictly at the disaster endpoint $s = 1$ with a margin exceeding this variation, it extends to the whole path; the pointwise and discrete characterisations can disagree only in a knife-edge band around $\tilde{\Lambda} = S$. Since $V_{mr}/V_{rr} > 0$ and $S_r \geq 0$, the reconstruction-side signal weakly raises $\tilde{\Lambda}$ and so makes defocusing more likely; it shuts down in the log-felicity limit, where $\tilde{\Lambda} = \Lambda$. Using the leading-order approximations $|\Delta u| \approx u'(AK_1)\delta AK_1$ and $(\Delta u)^2 \approx [u'(AK_1)]^2\delta^2 A^2 K_1^2$, the local comparative statics should be read as sign-contingent. Defocusing is favoured by high resource pressure—for example high disaster severity δ and high capital intensity K_0/K_1 , which raise Λ . When the mitigation-side diagnostic bracket

$$B_S \equiv \kappa_2 \mathcal{M} - \Theta \kappa_1^2 \mathcal{M}^2$$

is positive and the reconstruction-side signal S_r is small, lower political patience β ,⁵ lower mitigation effectiveness α , lower diagnostic intensity θ , and lower disaster-to-hazard transmission ψ_H weaken the positive diagnostic focusing channel and make defocusing more likely. When $B_S < 0$, by contrast, the belief shock reduces the marginal value of additional mitigation at the margin; in that region higher θ or higher ψ_H can reinforce defocusing rather than counteract it. Low probability-function curvature κ_2 and high slope κ_1 make this negative-signal case more likely by strengthening the saturation term relative to the mean term. Felicity curvature above the log case also pushes toward defocusing because $\Delta u_K \approx (\sigma - 1)\delta A u'(AK_1) > 0$ activates the reconstruction-side signal, raising $\tilde{\Lambda}$.

Risk sensitivity has a non-monotone effect. At $\Theta = 0$, $B_S = \kappa_2|\Delta u| > 0$ when $\kappa_2 > 0$. As Θ rises, the precautionary motive initially changes the value of probability reduction, but sufficiently high Θ can make the saturation term dominate. Whenever $\hat{p}_1\kappa_2 < \kappa_1^2$, the bracket turns negative beyond a finite threshold in Θ and approaches zero from below as $\Theta \rightarrow \infty$, where beliefs cease to matter at the margin. Past that threshold the diagnostic focusing channel itself works against mitigation, and any positive resource-defocusing channel favours defocusing.

A useful behavioral implication is immediate. Let $B_S \equiv \kappa_2\mathcal{M} - \Theta\kappa_1^2\mathcal{M}^2$. Holding the other local objects fixed, higher diagnosticity θ strengthens the diagnostic focusing channel when $B_S > 0$ and makes defocusing less likely. When $B_S < 0$, higher diagnosticity strengthens the negative saturation component and can make defocusing more likely (Appendix C).

This is the behavioral punchline of the model. More salience need not mean more mitigation. Diagnostic expectations always magnify the belief movement, but the policy consequence of that movement depends on whether the marginal value of additional probability reduction is still rising or has begun to saturate.

The comparative statics are therefore subtle because the diagnostic focusing channel has internal tension. Under empirically relevant parameters (a $p(\cdot)$ that is roughly logistic and convex in the lower tail, so $\kappa_2 > 0$ but moderate, with $\hat{p}_1$ small at baseline), the mean piece may dominate and $S > 0$. Defocusing then requires the resource-defocusing channel to exceed this positive signal. Under

⁵ Λ itself also scales with β , since $\mathcal{G}$ and $\mathcal{W}$ are proportional to β . From $V_{rr} = -(c + \mathcal{G})$ and $V_{mr} = -(c + \mathcal{W})$ with $c \equiv |u''(C_0)|$, one computes $\partial \ln \Lambda / \partial \ln \beta = c/(c + \mathcal{G}) \in (0, 1) < 1 = \partial \ln S / \partial \ln \beta$, so a fall in β lowers S proportionally faster than Λ and the stated direction survives β 's presence in the resource-side objects.

linear probability the mitigation-side signal is negative throughout, and defocusing follows from output destruction together with the negative saturation effect; there is no inversion, only a magnitude question.

Proposition 2 pins down the central qualitative prediction of the paper. Countries in the defocusing region cut mitigation in response to disasters; countries in the focusing region raise it. The 1990 Western European windstorm cluster sits comfortably on the defocusing side: treated countries combine modest political patience (typical of Western European institutional cycles), moderate diagnostic responsiveness, and high relative resource destruction ($\delta \sim 0.03$, matching the treated countries' damages of roughly 3% of GDP).

Before turning to the comparative statics of δ , it is useful to state explicitly what severity does and does not do in the model. In the baseline formulation, severity does not mechanically make the event more diagnostic; the belief shift is governed by D_0 , θ , and ψ_H . Severity instead scales the consequences of the event once it occurs. It raises the resource-defocusing term O , because the realised disaster destroys more period-0 output; it raises $|\Delta u|$, because the future disaster state becomes worse; and, when $\sigma > 1$, it raises Δu_K , strengthening the capital-buffer and reconstruction motive. The severity exercise therefore asks which of these severity-amplified effects grows faster.

Severity and the Defocusing Region

Severity does not decide the sign of the response at first order. A disaster twice as severe makes the conditional bad-state utility loss roughly twice as large, and it destroys twice the resources today: both channels scale one-for-one with δ at leading order, so doubling severity doubles the demand for mitigation and the squeeze on its supply in the same proportion. Whether a country defocuses is, to first order, a question about parameters rather than about how hard it was hit. What ties severity to defocusing is the second-order behaviour of the two channels, and there they bend in opposite directions.

The resource-defocusing channel compounds. A government that has just lost three per cent of output values each remaining unit of resources more than one that has lost one per cent, and not proportionally more: under CRRA the pain of giving up the next unit rises ever faster as consumption falls. This is prudence, $u''' > 0$, and it holds unconditionally. Each additional point of damage therefore raises the opportunity cost of funding mitigation by more than the last. Mitigation is the marginal casualty of this squeeze, not because voters value the climate less after a severe shock, but because the

shadow price of current resources rises convexly in damage. For $\sigma > 1$ a second force pushes the same way: fear of the next disaster makes holding capital into the future more valuable as a buffer, so the very belief shock that fuels demand for mitigation also recruits reconstruction as a rival claimant on the shrunken budget.

The diagnostic focusing channel saturates. Part of mitigation's appeal is insurance: it narrows the spread of possible futures, and a voter who prefers early resolution of uncertainty values that narrowing in itself. Insurance of this kind is worth most against tail risks that still seem remote. Once a shock raises perceived disaster probability, the voter's certainty equivalent already tilts toward the disaster state, and trimming the perceived chance a little buys correspondingly little reassurance; formally, the marginal value of probability reduction declines in perceived disaster probability at the rate $\Theta\mathcal{M}^2$, a global property of the certainty equivalent, and the saturation piece of S grows like δ^2 while the mean piece grows like δ . Prudence cuts the other way on the mean piece, since the utility loss $|\Delta u|$ is itself convex in damage; the signal saturates at second order when the curvature condition in Appendix B holds, which requires κ_2 to be moderate. Severe disasters then raise fear faster than they raise the usefulness of the instrument available to act on it.

This asymmetry, a compounding cost set against a saturating benefit, is the direct channel logic behind the upper-tail prediction. It should be read as a sufficient direct-effect comparison rather than an unconditional optimized comparative static: the full policy response also contains re-optimisation terms, which Appendix B does not sign under the maintained assumptions. This qualified prediction is the model's reading of Figure A.7: concern rises with damages everywhere, but legislation falls precisely where damages are largest.

A natural question is whether the asymmetry is an artefact of the binary shock, which fixes the belief revision at $(1 + \theta)\psi_H$ regardless of severity. Suppose instead that worse disasters are read as more diagnostic: $\psi_H(\delta) = \bar{\psi}_H \delta^\chi$ with $\chi \geq 0$, the baseline being $\chi = 0$. At the boundary of the defocusing region, where $\tilde{\Lambda} = S$, the region expands with severity if and only if the severity-elasticity of $\tilde{\Lambda}$ exceeds that of S . At leading order these elasticities, evaluated at the disaster endpoint $s = 1$ of the embedding with the local probability objects $(\hat{p}_1, \kappa_1, \kappa_2)$ held fixed at the evaluation point (a convention with content only when $\chi > 0$; see Appendix B), and assuming $\kappa_2 > 0$ and $(\tilde{b}/a)\delta < 1$, are (derivations

in Appendix B)

$$\varepsilon_S = 1 + \chi - \frac{(\tilde{b}/a)\delta}{1 - (\tilde{b}/a)\delta}, \quad \varepsilon_\Lambda = 1 + \omega(\sigma + 1) \frac{\delta AK_0}{C_0} - \tilde{w},$$

where a and $\tilde{b}$ are the linear and quadratic coefficients of the diagnostic bracket in δ (Appendix B), with $\tilde{b} > 0$, so that the signal saturates at second order, whenever $\Theta q \kappa_1^2 > \frac{\kappa_2}{2} [\sigma + \Theta q(1 - 2\hat{p}_1)]$ for $q \equiv AK_1 u'(AK_1)$; $\omega \equiv \beta A^2 u''(AK_1)/V_{rr} \in (0, 1)$; and

$$\tilde{w} \equiv \frac{\mathcal{W}}{\mathcal{G} - \mathcal{W}} = \frac{\Pi(\sigma - 1)}{\sigma - \Pi(\sigma - 1)} + O(\delta)$$

at an interior optimum is the severity drag from the belief–capital cross-wedge $\mathcal{W}$ inside V_{mr} : because $\mathcal{W} \propto \Delta u_K$ grows one-for-one with δ , it depresses ε_Λ by an amount of order $\sigma - 1$ rather than of order δ , and it vanishes identically in the log-felicity limit, where $\mathcal{W} \equiv 0$. For $\sigma > 1$ the reconstruction-side term adds a further share-weighted component to ε_Λ with elasticity $1 + \chi$ at leading order, which carries the same δ^χ scaling as S (Appendix B).

Discussion. Under severity-dependent diagnosticity $\psi_H(\delta) = \bar{\psi}_H \delta^\chi$ and the second-order saturation condition $\tilde{b} \geq 0$ of Appendix B, the direct-effect defocusing region expands with severity whenever $\varepsilon_\Lambda > \varepsilon_S$. At leading order, a sufficient condition for this direct-effect comparison is

$$\chi + \tilde{w} < \omega(\sigma + 1) \frac{\delta AK_0}{C_0},$$

which reduces to $\chi < \omega(\sigma + 1) \delta AK_0/C_0$ in the log-felicity limit, where $\tilde{w} = 0$.

Because $\tilde{w}$ is of order $\sigma - 1$ while the right-hand side is of order δ , the displayed condition binds only near log felicity: it requires $\sigma - 1$ comparable to δ , or severity large relative to the departure from log utility. In the log limit, severity-independent updating ($\chi = 0$) satisfies the condition trivially, though there the direct margin is itself of order δ and the unsigned re-optimisation terms discussed in Appendix B can matter quantitatively. On magnitude: at an interior optimum $\omega = (C_0/K_1)/(1 + C_0/K_1) \approx 1/2$ at leading order, so under log felicity with $\delta = 0.05$ and $AK_0/C_0 \approx 1.5$ the bound is roughly 0.075, rising linearly with severity; already at $\sigma = 1.05$ the wedge $\tilde{w} \approx 0.05$ absorbs two-thirds of that margin, and for σ appreciably above one the direct-effect comparison is quantitative rather than sign-determined. The condition is sufficient rather than necessary, with the saturation slack

widening the region whenever $\tilde{b} > 0$. For $\chi > 0$ the tail prediction survives when belief updating is sufficiently inelastic in severity relative to the convexity of the resource cost: what overturns it is not diagnosticity per se but a belief response more convex in severity than the prudence-driven resource cost. Diminishing sensitivity to the magnitude of news, the natural case under the salience foundations of diagnostic expectations (N. Gennaioli and Shleifer, 2010; Bordalo, N. Gennaioli, and Shleifer, 2018), pushes toward small χ ; how small is ultimately a quantitative question, and the descriptive evidence in Figure A.7, where concern rises with damages while legislation falls, indicates that the resource side dominates observationally in the relevant range.

3.3 Institutional Design Implications

The defocusing condition also gives direct policy-design comparative statics. Institutions matter because they shift the relative size of the resource-defocusing channel and the diagnostic focusing channel.

The model delivers a direct institutional-design implication. Any institution that lowers the post-disaster resource cost O , lowers the reconstruction-side crowding-out term S_r , or strengthens the positive part of the diagnostic focusing channel S shrinks the defocusing region. Conversely, institutions that raise the shadow cost of current resources or shorten political horizons expand it (Appendix C).

Pre-arranged disaster finance, catastrophe insurance, and automatic stabilisers reduce the effective resource squeeze after a disaster and therefore lower O . Ring-fenced climate budgets or green-reconstruction rules limit the extent to which reconstruction competes with mitigation, reducing the adjusted output term $\tilde{\Lambda}$. Institutions that lengthen political horizons raise the effective continuation weight β and strengthen the positive diagnostic focusing channel when $B_S > 0$. Finally, policies that make reconstruction complementary with mitigation—for example by requiring rebuilding to use low-carbon capital—raise the productive co-benefit of mitigation in the dynamic model and weaken the persistence of the mitigation shortfall.

3.4 A Cognitive Counterpart to Fiscal Constraints

The resource constraint in the model bites through fiscal capacity, but the same logic has a natural cognitive reading, which we sketch informally. Bordalo, N. Gennaioli, Lanzani, et al. (2025) develop a theory

of selective attention in which decision-makers categorise problems by similarity to past experiences and allocate their attention budget to the features most salient in the active category. A disaster shock plausibly triggers a “crisis mode” schema, with immediate output losses, reconstruction, and emergency governance as the goal-relevant features, rather than a “planning mode” schema in which long-run climate hazard reduction is central. The shift is self-reinforcing. The salience that drives diagnostic overweighting (our diagnostic focusing channel) is the same salience that cues a reconstruction-focused attentional frame, which then crowds out the legislative and administrative bandwidth needed to produce mitigation policy.

The cognitive account complements the fiscal one. Even with slack budgets, attentional crowd-out alone would generate a wedge between heightened perceived disaster probability and observed mitigation supply. The two channels reinforce each other: fiscal constraints limit the resources available, and attentional crowd-out limits the political will and administrative capacity to deploy whatever resources remain. Both push in the same direction on Theorem 2, widening the gap between the signal and resource-defocusing effects and making the translation failure both more likely and more persistent.

3.5 Dynamic Model

The two-period model establishes when a disaster reduces mitigation on impact. To ask how long this effect lasts, we embed the same mechanism in an infinite-horizon setting. The purpose of this section is not to obtain a global closed-form solution. It is to show, locally around the non-disaster steady state, how a temporary shock can generate persistent mitigation shortfalls. Persistence is governed by the eigenvalues of the linearised capital-hazard system: when the capital/resource-defocusing component is more persistent than the belief component, the resource channel can remain active long after diagnostic beliefs have reverted.

Infinite-Horizon Model

Time is discrete, $t = 0, 1, 2, \dots$. The incumbent solves

$$\max_{\{m_t, r_t\}} V_0,$$

where V_0 is given recursively by the risk-sensitive aggregator (7) with diagnostic beliefs, subject to

$$C_t = AK_t(1 - \delta D_t) - m_t - r_t \quad (23)$$

$$K_{t+1} = (1 - d)K_t + r_t + \xi m_t \quad (24)$$

$$H_{t+1} = \rho H_t - \alpha m_t + \psi_H D_t + \eta_t \quad (25)$$

with $\xi \in [0, 1]$ a productive co-benefit of mitigation through investment in clean capital, and $\eta_t \sim N(0, \sigma_\eta^2)$ i.i.d. climate variability. The Bellman equation is the dynamic counterpart of (7),

$$V(K_t, H_t) = \max_{m_t, r_t} \{u(C_t) + \beta \mathcal{R}_t[V(K_{t+1}, H_{t+1})]\}, \quad \mathcal{R}_t[X] = -\frac{1}{\Theta} \ln \mathbb{E}_t^{DE}[e^{-\Theta X}],$$

and the impulse responses below are computed by linearising the recursion around the local steady-state restrictions stated below. The expectation operator $\mathbb{E}_t^{DE}$ over the joint disaster–Gaussian innovation, the perceived-hazard-state recursion it induces, and the MIT-shock convention behind the impulse responses are formalised in Appendix D.

Steady State, Local Linearisation, and Impulse Responses

We analyse the model’s dynamics by linearising around the non-disaster steady state. The impulse-response result below is therefore a local persistence result. It should be read as a disciplined way to decompose propagation into a belief component and a capital/resource-defocusing component, rather than as a global comparative-static theorem.

Steady-state restrictions. Let $V^{*,n}$ and $V^{*,d}$ denote the steady-state continuation values conditional on a no-disaster and a disaster realisation next period, $\Delta V^* \equiv V^{*,n} - V^{*,d} > 0$ the continuation-value loss from a disaster, and

$$\mathcal{M}_V^* \equiv \frac{e^{\Theta \Delta V^*} - 1}{\Theta[1 + p^*(e^{\Theta \Delta V^*} - 1)]}$$

the marginal value of probability reduction at the value level, with $p^* \equiv p(H^*)$. The non-stochastic steady state ($D_t = 0 \forall t$) satisfies the following local first-order restrictions

$$u'(C^*) = \beta V_K^* \quad (26)$$

$$V_K^* = \frac{u'(C^*)A}{1 - \beta(1 - d)} \quad (27)$$

$$V_H^* = -\frac{p'(H^*)\mathcal{M}_V^*}{1 - \beta\rho}, \quad (28)$$

where (26)–(27) hold up to second-order corrections of order $p^*\delta$ (the tilted-mean adjustment across the period-ahead disaster states) and $\Theta\sigma_\eta^2$ (the Gaussian tilt), both absorbed by the linearisation, and (28) holds to the same order: given ΔV^* , its derivation replaces the two branch-specific continuation derivatives by a common value, and the resulting error is the $O(\delta)$ cross-branch variation in V_H' identified in Appendix D, likewise absorbed by the linearisation. By the envelope theorem, $\Delta V^* = \delta AK^*u'(C^*) + \beta(1 + \theta)\psi_H|V_H^*| + O(\delta^2 + \psi_H^2\delta)$; as $\Theta \rightarrow 0$, $\mathcal{M}_V^* \rightarrow \Delta V^*$, so (28) nests the expected-utility form $V_H^* = -p'(H^*)\delta AK^*u'(C^*)/(1 - \beta\rho)$ at leading order in δ and ψ_H . well posed with a positive solution provided $\lambda\Pi_V^* < 1$ at the solution, where $\lambda \equiv \beta(1 + \theta)\psi_H p'(H^*)/(1 - \beta\rho)$ and $\Pi_V \equiv \partial\mathcal{M}_V/\partial\Delta V = e^{\Theta\Delta V}/[1 + p^*(e^{\Theta\Delta V} - 1)]^2$; since $\Delta V^* = O(\delta)$, the condition reduces to $\lambda < 1$ at leading order and exactly as $\Theta \rightarrow 0$, which we assume.

Existence, the mitigation margin, and stability. Because production is linear in capital, (26)–(27) jointly imply $\beta(A + 1 - d) = 1$ up to the stated $O(p^*\delta + \Theta\sigma_\eta^2)$ corrections: with AK technology a deterministic interior steady state exists only on a neighbourhood of this knife-edge, the disaster-risk and Gaussian-tilt wedges supplying the slack that pins (K^*, C^*) . We therefore read $\beta(A + 1 - d) - 1 = O(p^*\delta + \Theta\sigma_\eta^2)$ as a maintained restriction on primitives rather than a free parameter combination. The steady state must in addition satisfy the first-order condition for mitigation, $u'(C^*) = \beta\xi V_K^* + \alpha\beta|V_H^*|$ up to the same corrections, equivalently $(1 - \xi)V_K^* = \alpha|V_H^*|$, which ties the hazard block to the capital block and is required for $m^* > 0$; it is the dynamic counterpart of (FOC-m). Finally, the stability assumption in Proposition 3 is substantive in this environment: the open-loop capital root is $1 - d + h_K + \xi g_K$, and with AK technology mean reversion in the capital block comes entirely from the policy feedbacks (h_K, g_K) , so stability of A_x should be verified in any calibration rather than presumed.

Now consider an unanticipated disaster at $t = 0$: $D_0 = 1$, $D_t = 0$ for $t \geq 1$.

Proposition 3. *Suppose the dynamic problem is differentiable around the non-disaster steady state and the linearised system is stable. Let deviations from steady state be denoted by tildes and let the local policy rules be*

$$\tilde{m}_t = g_K \tilde{K}_t + g_H \tilde{H}_t + g_D D_t, \quad \tilde{r}_t = h_K \tilde{K}_t + h_H \tilde{H}_t + h_D D_t.$$

Define $x_t \equiv (\tilde{K}_t, \tilde{H}_t)'$. Consider a one-time unanticipated disaster shock: $D_0 = 1$ and $D_t = 0$ for all $t \geq 1$. The linearised state system is

$$x_{t+1} = A_x x_t + b_D D_t,$$

where

$$A_x = \begin{pmatrix} 1 - d + h_K + \xi g_K & h_H + \xi g_H \\ -\alpha g_K & \rho - \alpha g_H \end{pmatrix}, \quad b_D = \begin{pmatrix} h_D + \xi g_D \\ (1 + \theta)\psi_H - \alpha g_D \end{pmatrix}.$$

The mitigation impulse response is

$$IRF_m(0) = g_D, \quad IRF_m(t) = \begin{pmatrix} g_K & g_H \end{pmatrix} A_x^{t-1} b_D \quad \text{for } t \geq 1.$$

The cumulative response over T periods is therefore

$$\sum_{t=0}^T IRF_m(t) = g_D + \begin{pmatrix} g_K & g_H \end{pmatrix} \left(\sum_{j=0}^{T-1} A_x^j \right) b_D,$$

and, if $I - A_x$ is invertible,

$$\sum_{t=0}^T IRF_m(t) = g_D + \begin{pmatrix} g_K & g_H \end{pmatrix} (I - A_x^T)(I - A_x)^{-1} b_D.$$

A simple two-mode representation makes the persistence logic transparent. Suppose the local impulse response can be represented by two dominant modes,

$$IRF_m(t) = a\lambda_H^t - b\lambda_K^t,$$

where $a > 0$ is the diagnostic focusing component, $b > 0$ is the resource-defocusing component, $\lambda_H \in (0, 1)$ is the belief-persistence eigenvalue, and $\lambda_K \in (0, 1)$ is the capital/resource eigenvalue; at $t = 0$ the representation identifies $g_D = a - b$ in Proposition 3. If $b > a$, mitigation falls on impact, and if in addition $\lambda_K \geq \lambda_H$ the response is negative at every horizon: defocusing is immediate and, within the local approximation, permanent. If instead $a > b$ and $\lambda_K > \lambda_H$, the belief mode dominates initially and the response turns negative at the crossing date $t^* = \ln(a/b)/\ln(\lambda_K/\lambda_H)$, after which it remains negative; call the dates with $IRF_m(t) < 0$, here $[t^*, \infty)$, the defocusing spell. A larger λ_K/λ_H brings the onset forward, since t^* is decreasing in λ_K/λ_H ; deepens the shortfall at every subsequent date, since $\partial IRF_m(t)/\partial \lambda_K = -bt\lambda_K^{t-1} < 0$; and lowers the cumulative response $\sum_{t \geq 0} IRF_m(t) = a/(1 - \lambda_H) - b/(1 - \lambda_K)$. In this sense the depth and duration of defocusing increase with λ_K/λ_H even when the initial belief response is strong.

No embedding is needed here. Within the linearisation, the policy rules are linear in the shock vector, so the response to the binary disaster coincides exactly with the difference between the $D_0 = 1$ and $D_0 = 0$ paths; the notation $IRF_m(t)$ denotes this difference. The representation separates two propagation mechanisms. The diagnostic focusing component enters through the perceived-hazard-state element of b_D , which contains the impact belief shift $(1 + \theta)\psi_H$. The resource-defocusing component enters through the direct policy response coefficients (g_D, h_D) , which capture the contemporaneous resource loss in (23), and through the capital equation in (24). The output-propagation rate is generally the dominant capital/output eigenvalue of A_x , not exactly $1 - d$; it is approximately $1 - d$ only when policy feedbacks through g_K, g_H, h_K, h_H and the clean-capital co-benefit ξ are small. Under calibrations in which the dominant capital/output eigenvalue is larger than the belief-persistence eigenvalue, the induced capital shortfall outlasts the spike in attention. This gives the dynamic version of the translation failure: the disaster can focus beliefs on climate risk immediately, while the policy system remains defocused by the slower-moving resource constraint.

The sign of a belief shock is also governed by the full two-instrument response, not only by the mitigation-side bracket. Locally, the pure belief effect has the same structure as Theorem 2: its numerator contains


$$-V_{rr} \beta \alpha (\kappa_2 \mathcal{M} - \Theta \kappa_1^2 \mathcal{M}^2) + V_{mr} \beta \kappa_1 \Pi \Delta u_K.$$

The second term is weakly negative because $V_{mr} < 0$ and $\Delta u_K \geq 0$. In the log-felicity case, or whenever

the reconstruction-side signal is negligible, this reduces to the sign of the mitigation-side bracket in (18).

Figure 2 illustrates the decomposition for a representative country in the defocusing parameter region. The diagnostic focusing contribution peaks on impact and decays with the belief-persistence eigenvalue. The resource-defocusing contribution is more persistent when the dominant capital/output eigenvalue is larger, so resource constraints remain materially tighter after belief distortions have mostly reverted. The net response is negative on impact and can deepen as the diagnostic focusing channel decays faster than the resource-defocusing channel. Mitigation reaches its trough roughly eight years after the disaster and remains below baseline well beyond a decade in the calibration shown. The figure should therefore be read as an illustration of the local eigenvalue mechanism, not as an implication of a fixed depreciation-only decay rate.

Figure 2: Calibrated impulse response of mitigation to a one-time disaster



Notes: Linearised impulse response of mitigation m_t to a one-time disaster $D_0 = 1$ in the infinite-horizon model under risk-sensitive preferences, for a country in the defocusing parameter region. The solid line plots the total response; the dashed line plots the diagnostic focusing contribution, and the dotted line plots the resource-defocusing contribution through capital dynamics. In the calibration the capital/output eigenvalue is close to $1 - d$. The shaded region marks the cumulative mitigation shortfall.

4 Conclusion

This paper has developed a diagnostic-expectations theory of climate focusing events. The central message is that disasters can focus attention while defocusing policy. A realised disaster is diagnostic evidence about a latent hazard state, so it raises the perceived probability of future climate losses. Yet the same disaster destroys current resources and creates reconstruction needs. Mitigation falls when the resource-defocusing channel dominates the diagnostic focusing channel.

Disaster severity is central to this trade-off. More severe shocks make future losses larger and can increase the value of mitigation, but they also tighten the current resource constraint and make reconstruction more urgent. The model therefore does not predict a monotone relationship between severity and mitigation. Its upper-tail implication is that very severe shocks can be especially prone to defocusing when the resource squeeze compounds while the marginal value of further probability reduction saturates.

The formal analysis is organised around this wedge. Proposition 1 shows that disasters are focusing events in beliefs: perceived hazard and perceived disaster probability rise, and diagnosticity amplifies the response. Theorem 2 decomposes the mitigation response into diagnostic focusing, resource defocusing, and reconstruction-side signal terms. Proposition 2 characterises the attention–policy translation failure: mitigation falls if and only if the adjusted resource channel exceeds the diagnostic focusing channel along the exposure path. The comparative statics highlight the behavioral implication that stronger diagnosticity need not always raise mitigation, because the probability-reduction value of mitigation can saturate. The dynamic extension shows how the policy shortfall can persist when the capital/resource eigenmode decays more slowly than the belief eigenmode.

The empirical section is intended only as motivation. The 1990 Western European windstorm cluster illustrates the pattern that the model explains: large shocks can coincide with rising concern, slower mitigation legislation, and tighter resources. The theory does not depend on that episode as a definitive causal test. Instead, the episode serves to discipline the central question: why might a society become more worried about climate risk while producing less mitigation policy?

The institutional implication is that translating disaster salience into mitigation requires capacity, not just attention. Pre-arranged disaster financing, automatic stabilisers, ring-fenced mitigation budgets, and green-reconstruction rules all reduce the resource-defocusing channel. Longer political horizons and

stronger commitment devices raise the effective continuation value of mitigation. These institutions matter precisely because the model does not deny that disasters focus attention; it shows why attention alone is insufficient.

Several extensions are natural. The theory could incorporate heterogeneous voter groups, partisan competition, endogenous legislative capacity, or a planner benchmark. Empirically, the most useful next step is to measure the two margins directly (beliefs and capacity) using parliamentary speech, media text, emergency spending, reconstruction budgets, and environmental-agency resources. Such evidence would test whether the diagnostic focusing channel and the resource-defocusing channel move as the model predicts.

References

- Afrouzi, Hassan et al. (2023). “Overreaction in expectations: Evidence and theory”. In: *The Quarterly Journal of Economics* 138.3, pp. 1713–1764.
- Andre, Peter et al. (2024). “Globally Representative Evidence on the Actual and Perceived Support for Climate Action”. In: *Nature Climate Change* 14.3, pp. 253–259. ISSN: 1758-6798. DOI: 10.1038/s41558-024-01925-3. (Visited on 03/12/2024).
- Bansal, Ravi and Amir Yaron (2004). “Risks for the long run: A potential resolution of asset pricing puzzles”. In: *The journal of Finance* 59.4, pp. 1481–1509.
- Barrot, Jean-Noël and Julien Sauvagnat (2016). “Input specificity and the propagation of idiosyncratic shocks in production networks”. In: *Quarterly Journal of Economics* 131.3, pp. 1543–1592.
- Besley, Timothy and Torsten Persson (2009). “The Origins of State Capacity: Property Rights, Taxation, and Politics”. In: *American Economic Review* 99.4, pp. 1218–44. DOI: 10.1257/aer.99.4.1218. URL: <https://www.aeaweb.org/articles?id=10.1257/aer.99.4.1218>.
- (2011). *Pillars of prosperity: The political economics of development clusters*. Princeton University Press.
- Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo (2024). “Diagnostic Business Cycles”. In: *The Review of Economic Studies* 91.1, pp. 129–162.
- Birkland, Thomas A. (1998). “Focusing Events, Mobilization, and Agenda Setting”. In: *Journal of Public Policy* 18.1, pp. 53–74. ISSN: 1469-7815, 0143-814X. DOI: 10.1017/S0143814X98000038. (Visited on 04/04/2024).
- (2006). *Lessons of disaster: Policy change after catastrophic events*. Georgetown University Press.
- Bordalo, Pedro, Nicola Gennaioli, Giacomo Lanzani, et al. (2025). *A cognitive theory of reasoning and choice*. Tech. rep. National Bureau of Economic Research.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018). “Diagnostic Expectations and Credit Cycles”. In: *Journal of Finance* 73.1, pp. 199–227. DOI: 10.1111/jofi.12586.

- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2022). “Overreaction and diagnostic expectations in macroeconomics”. In: *Journal of Economic Perspectives* 36.3, pp. 223–244.
- Carattini, Stefano et al. (2023). “What does network analysis teach us about international environmental cooperation?” In: *Ecological Economics* 205, p. 107670.
- Cohen, Charles and Eric D Werker (2008). “The political economy of “natural” disasters”. In: *Journal of Conflict Resolution* 52.6, pp. 795–819.
- de Silva, Tiloka and Silvana Tenreyro (Dec. 2021). “Presidential Address 2021 Climate-Change Pledges, Actions, and Outcomes”. In: *Journal of the European Economic Association* 19.6, pp. 2958–2991. ISSN: 1542-4766. DOI: 10.1093/jeea/jvab046. (Visited on 04/22/2022).
- Dixit, Avinash and John Londregan (1996). “The determinants of success of special interests in redistributive politics”. In: *the Journal of Politics* 58.4, pp. 1132–1155.
- Epstein, Larry G. and Stanley E. Zin (1989). “Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework”. In: *Econometrica* 57.4, pp. 937–969.
- Eskander, Shaikh MSU and Sam Fankhauser (2023). “The Impact of Climate Legislation on Trade-Related Carbon Emissions 1996–2018”. In: *Environment and Resource Economics* 85, pp. 167–194. DOI: <https://doi.org/10.1007/s10640-023-00762-w>.
- Fankhauser, Samuel, Caterina Gennaioli, and Murray Collins (Apr. 2016). “Do International Factors Influence the Passage of Climate Change Legislation?” In: *Climate Policy* 16.3, pp. 318–331. ISSN: 1469-3062, 1752-7457. DOI: 10.1080/14693062.2014.1000814. (Visited on 10/05/2022).
- Galanis, Giorgos, Giorgio Ricchiuti, and Ben Tippet (2025). *Heterogeneity and Global Climate Action*. Tech. rep. Grantham Research Institute on Climate Change and the Environment Working Paper No. 436.
- Garrett, Thomas A and Russell S Sobel (2003). “The political economy of FEMA disaster payments”. In: *Economic inquiry* 41.3, pp. 496–509.
- Gaspar, John T and Andrew Reeves (2011). “Make it rain? Retrospection and the attentive electorate in the context of natural disasters”. In: *American journal of political science* 55.2, pp. 340–355.

- Gennaioli, Nicola and Andrei Shleifer (2010). “What Comes to Mind”. In: *Quarterly Journal of Economics* 125.4, pp. 1399–1433. DOI: 10.1162/qjec.2010.125.4.1399.
- Giordono, Leanne, Hilary Boudet, and Alexander Gard-Murray (Dec. 2020). “Local Adaptation Policy Responses to Extreme Weather Events”. In: *Policy Sciences* 53.4, pp. 609–636. ISSN: 1573-0891. DOI: 10.1007/s11077-020-09401-3. (Visited on 04/04/2024).
- Hansen, James, Makiko Sato, and Reto Ruedy (2012). “Perception of climate change”. In: *Proceedings of the National Academy of Sciences* 109.37, E2415–E2423. DOI: 10.1073/pnas.1205276109.
- Hansen, Lars Peter, John C. Heaton, and Nan Li (2008). “Consumption strikes back? Measuring long-run risk”. In: *Journal of Political Economy* 116.2, pp. 260–302.
- Hansen, Lars Peter and Thomas J Sargent (2008). *Robustness*. Princeton University Press.
- Healy, Andrew and Neil Malhotra (2009). “Myopic voters and natural disaster policy”. In: *American Political Science Review* 103.3, pp. 387–406.
- Herrnstadt, Evan and Erich Muehlegger (2014). “Weather, Salience of Climate Change and Congressional Voting”. In: *Journal of Environmental Economics and Management* 68.3, pp. 435–448. DOI: 10.1016/j.jeem.2014.08.002.
- Hoffmann, Roman et al. (2022). “Climate Change Experiences Raise Environmental Concerns and Promote Green Voting”. In: *Nature Climate Change* 12.2, pp. 148–155. ISSN: 1758-6798. DOI: 10.1038/s41558-021-01263-8. (Visited on 04/04/2024).
- IPCC (2022). *Climate Change 2022 – Impacts, Adaptation and Vulnerability: Working Group II Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. H.-O. Pörtner, D.C. Roberts, M. Tignor, E.S. Poloczanska, K. Mintenbeck, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, B. Rama (Eds.) Cambridge, UK and New York, NY: Cambridge University Press. ISBN: 978-1-00-932584-4. (Visited on 03/13/2024).
- Kahn, Matthew E (2005). “The death toll from natural disasters: the role of income, geography, and institutions”. In: *Review of economics and statistics* 87.2, pp. 271–284.

- Koks, E. E. and T. Haer (2020). “A High-Resolution Wind Damage Model for Europe”. In: *Scientific Reports* 10.1, p. 6866. ISSN: 2045-2322. DOI: 10.1038/s41598-020-63580-w. (Visited on 10/10/2024).
- L’Huillier, Jean-Paul, Sanjay R Singh, and Donghoon Yoo (2024). “Incorporating Diagnostic Expectations into the New Keynesian Framework”. In: *The Review of Economic Studies* 91.5, pp. 3013–3046.
- Lindbeck, Assar and Jörgen W Weibull (1987). “Balanced-budget redistribution as the outcome of political competition”. In: *Public choice* 52.3, pp. 273–297.
- Lohmann, Paul M. and Andreas Kontoleon (2023). “Do Flood and Heatwave Experiences Shape Climate Opinion? Causal Evidence from Flooding and Heatwaves in England and Wales”. In: *Environmental and Resource Economics* 86.1–2, pp. 263–304. DOI: 10.1007/s10640-023-00796-0.
- Osberghaus, Daniel and Carina Fugger (2022). “Natural Disasters and Climate Change Beliefs: The Role of Distance and Prior Beliefs”. In: *Global Environmental Change* 74, p. 102515. DOI: 10.1016/j.gloenvcha.2022.102515.
- Otto, Friederike E. L. et al. (2021). “Exposure to natural hazard events unassociated with policy change for improved disaster risk reduction”. In: *Nature Climate Change* 11, pp. 112–119. DOI: 10.1038/s41558-021-01222-3.
- Rahmstorf, Stefan and Dim Coumou (2011). “Increase of Extreme Events in a Warming World”. In: *Proceedings of the National Academy of Sciences* 108.44, pp. 17905–17909. DOI: 10.1073/pnas.1101766108. (Visited on 12/26/2024).
- Restoy, Fernando and Philippe Weil (2011). “Approximate equilibrium asset prices”. In: *Review of Finance* 15.1, pp. 1–28.
- Rowan, Sam (2022). “Extreme Weather and Climate Policy”. In: *Environmental Politics*, pp. 1–24. ISSN: 0964-4016, 1743-8934. DOI: 10.1080/09644016.2022.2127478. (Visited on 04/03/2024).
- Sauquet, Alexandre (2014). “Exploring the nature of inter-country interactions in the process of ratifying international environmental agreements: the case of the Kyoto Protocol”. In: *Public Choice* 159.1, pp. 141–158.
- Sloggy, Matthew R. et al. (2021). “Changing Climate, Changing Minds? The Effects of Natural Disasters on Public Perceptions of Climate Change”. In: *Climatic Change* 168.3–4, p. 25. DOI: 10.1007/s10584-021-03242-6.

- Swiss Re Institute (2024). *Natural catastrophes and man-made disasters in 2023: Gearing up for today's and tomorrow's weather risks*. Report sigma No. 1/2024. Swiss Re Institute. URL: <https://www.swissre.com/institute/research/sigma-research/sigma-2024-01.html> (visited on 03/15/2024).
- Tallarini Jr, Thomas D (2000). "Risk-sensitive real business cycles". In: *Journal of Monetary Economics* 45.3, pp. 507–532.
- Wappenhans, Tim et al. (2024). "Extreme Weather Events Do Not Increase Political Parties' Environmental Attention". In: *Nature Climate Change* 14.7, pp. 696–699. ISSN: 1758-6798. DOI: 10.1038/s41558-024-02024-z. (Visited on 12/26/2024).
- Weil, Philippe (1990). "Non-expected utility in macroeconomics". In: *Quarterly Journal of Economics* 105.1, pp. 29–42.

A Empirical Illustration

This section provides a brief empirical illustration of the paper’s mechanism. Our aim is to document a pattern consistent with the model’s *defocusing* channel: large disaster shocks coincide with (i) a temporary slowdown in mitigation policy production and (ii) a tightening of resources.

A.1 Data Sources and Sample

We measure climate-related disasters using the Emergency Events Database (EM-DAT) and follow the IMF classification, defining climate-related disasters as droughts, extreme temperatures, floods, landslides, storms, and wildfires. We use annual country-level counts. EM-DAT records events that meet at least one of the following criteria: ten or more deaths, 100 or more affected/injured, a declared state of emergency, or a request for international assistance.

Domestic mitigation legislation is taken from the Climate Change Laws of the World database. Our baseline outcome focuses on mitigation-only laws (excluding laws primarily classified as adaptation or disaster response). We focus on this narrow definition in order to try to capture a relatively comparable measure of mitigation action. One concern with this indicator may be that not all mitigation laws are equally effective at actually reducing emissions. Previous papers however have shown that mitigation laws captured in this dataset do effectively reduce GHG emissions de Silva and Tenreyro, 2021. Moreover, Table ?? shows in a simple fixed effects model that implementing a mitigation law in a given period is significantly associated with a reduction in GHG emissions contemporaneously and in the next period.

Table A.1: The impact of mitigation laws on total GHG emissions

	(1) GHG b/p
Mitigation laws (t)	-10.361** (0.047)
Mitigation laws (t-1)	-9.461* (0.060)
Mitigation laws (t-2)	-9.915 (0.153)
N	624
P-value in parentheses * 0.10 **0.05 ***0.01 Country and Year fixed effects are included	

We use direct economic damages to GDP from these extreme weather events to measure the extent to which a disaster shock led to a tightening of resources. We use direct economic damages to GDP (%) rather than aggregate statistics (changes in the level or trend of GDP, the capital stock, or government spending), as aggregate indicators will be determined by broad macroeconomic factors, not just extreme

weather events. Direct damages to GDP from extreme weather events are drawn from the EM-DAT database and cover the same extreme weather events discussed above.

Our sample consists of Western European countries from 1980–2020.⁶ We focus on this region because countries share broadly similar institutional and economic environments and because climate policy was politically salient over the period.⁷

A.2 Treatment Definition: The 1990 Windstorm Cluster

Motivated by the model, we focus on unusually large disasters relative to a country’s typical experience. Our baseline definition classifies an *extreme* year as one in which the annual count of climate-related disasters exceeds the country mean by four standard deviations. This threshold follows common definitions of extremes in the climate-statistics literature (J. Hansen, Sato, and Ruedy, 2012; Rahmstorf and Coumou, 2011).⁸ Annual counts of climate-related disasters, rather than damages to GDP, are used to identify the exogenous shock, as damages to GDP from natural disasters are determined by endogenous factors, such as the history of adaptation and risk management investments. On the other hand, an abnormal number of extreme events hitting a country in a given year is highly exogenous and unpredictable.

A surprising result of this data-driven identification strategy is that almost all the shocks occur in 1990.⁹ The year 1990 was particularly bad for climate-related disasters in Western Europe. Eight major storms—Daria, Herta, Judith, Nana, Ottilie, Polly, Vivian, and Wiebke—hit the continent, causing historically high levels of damage.¹⁰ These events represent a highly unusual concentration of extreme weather activity, with multiple countries experiencing disaster counts that are more than four standard deviations above their historical average.

The treatment group consists of countries crossing the extreme threshold in 1990 (Finland, Netherlands, Luxembourg, and Denmark); the remaining Western European countries form the control group.

These shocks were concentrated within 35 days¹¹ hence providing a natural experiment setting that is both temporally sharp and plausibly exogenous, satisfying the conditions for a valid event study. As climate disasters of this magnitude are unpredictable and outside the scope of normal political planning,

⁶We define Western Europe geographically to include Austria, Belgium, Denmark, Finland, France, Germany, Ireland, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland and the United Kingdom.

⁷The IPCC was established in 1988, while *Time Magazine* put the endangered earth on their front cover at the beginning of 1989, indicating that it had partly seeped into the cultural consciousness.

⁸Appendix A.5 considers alternative cutoffs and related definitions.

⁹The only exception is Spain in 2019, which we ignore as it is at the end of the sample and therefore its effects cannot be analysed.

¹⁰See Koks and Haer, 2020, which uses OpenStreetMap data to confirm the exceptional nature of damages in 1990. To our knowledge this is the only other paper that highlights this year as a significant confluence of events.

¹¹See Table A.6, in the Appendix A.6, for details.

we treat the 1990 shock as an exogenous treatment affecting only the identified countries. Importantly, 1990 also precedes major international climate agreements (e.g., the 1992 UNFCCC and 1997 Kyoto Protocol), making it a natural focal point for evaluating how extreme events may have influenced the domestic climate policy trajectory during a critical formative period.

A.3 Difference-in-Differences Specification

Following the predictions of the theoretical model, we estimate the impact of an extreme disaster shock on two outcome variables.

First, we estimate the average treatment effect by comparing the cumulative number of mitigation laws passed in the treated countries to those in the control group of Western European countries that did not experience such an extreme shock.

Second, we estimate the impact of the 1990 shock on realised GDP damages, to assess whether the impact on mitigation policies is likely to be driven by the mechanisms suggested by the model. We compute the average treatment effect by comparing the GDP damages in treated countries with those in a control group of Western European countries that did not experience a comparable extreme event. This approach allows us to assess whether the 1990 shock did constrain economic resources in the affected countries.

While it would also be useful to conduct a similar diff-in-diff regression with the measure of public concern over climate change, there is not sufficient data to do this. However, Figure A.7 above does provide some preliminary evidence that countries which experienced the highest damages to GDP did implement fewer laws—a result that is consistent with the story of this paper.

Given the common shock year, we use a simple difference-in-differences specification to summarize the post-1990 divergence in mitigation policy output:

$$y_{it} = \beta^{DD} (Treated_i \times Post_t) + \Gamma' X_{it} + \mu_i + \tau_t + \varepsilon_{it}, \quad (\text{A.1})$$


where y_{it} is either (i) the cumulative number of mitigation policies since 1980 or (ii) the damages to GDP of extreme weather events since 1980, μ_i and τ_t are country and year fixed effects, and X_{it} includes standard covariates (fossil fuel rents, GDP per capita, population, democracy, and government ideology). We report results over 1980–2000 to focus on medium-run responses around the episode; Appendix A.5 considers alternative horizons and specifications.

The figure we present on mitigation laws use the cumulative numbers to summarize medium-run output. Alternative results using annual counts and count models are reported in Appendix A.5.

A.4 Baseline Results

Figure A.1 presents the change in the cumulative number of mitigation laws between the treated and control groups. As discussed above, the climate disaster shock occurs in 1990 which is indicated by the

Figure A.1: Impact of climate disaster shock on cumulative no. mitigation laws



Note: 1990 is identified as the only year where the number of climate damages 4σ above the country mean. Treated group = Finland, Netherlands, Luxembourg and Denmark. Control group = Austria, Belgium, France, Germany Ireland, Italy, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom.

dark vertical line. The parallel trends assumption holds for the pre-treatment period, as both groups have the same flat trends. In 1990, the control group starts to implement mitigation laws, while the treated group does not implement laws until 1997¹². The parallel trends assumption is that both groups would have followed the same trajectory had the treated group not suffered from the disaster shock in 1990.

Table A.2 columns 1 and 2 reports the estimated post-1990 divergence in mitigation policy output between treated and control countries using equation (A.1). Across specifications, the interaction term $Treated_i \times Post_t$ is negative, indicating a relative slowdown in mitigation policy production in the treated group following the 1990 disaster episode. Each treated country had on average 0.303 fewer mitigation laws than the control countries in the post shock period from 1990 to 2000, significant at the 5% level. This effect is the same, in terms of magnitude, as the average number of mitigation laws of each country in this period, which was 0.30 cumulative laws per country.

The negative association is robust to including standard covariates (column 2), although none of

¹²The increase in mitigation laws for both the treatment and the control group in 1997 is driven by all countries signing the 1997 Kyoto Accords

Table A.2: The impact of climate disaster shock on mitigation laws and damages to GDP

	(1) Laws	(2) Laws	(3) Damage/GDP
Interaction	-0.303** (0.026)	-0.338** (0.045)	0.119* (0.090)
Controls			
Fossil fuel rents/GDP		0.009 (0.693)	
Population		-0.000 (0.563)	
Left-right index		-0.028 (0.482)	
GDP per capita		-0.006 (0.540)	
Democracy index		-3.358 (0.439)	
N	336	315	336

P-value in parentheses | * 0.10 **0.05 ***0.01 | Country + Year fixed effects included

the covariates are significant. We interpret these patterns as consistent with the first aspect of the model’s defocusing channel: even if disasters raise perceived risk (as shown in figure A.7a), the supply of mitigation policy can actually decline.

We next investigate the extent to which this decline in mitigation laws was driven by a constraint on resources. Figure A.2 presents the damages to GDP between the treated and control groups. In 1990, there is a clear upward shift in damages to GDP (and therefore decline in GDP) in the treated countries, without any significant change in the control group, which continues on its pre-treatment trend. Table A.2 column 3 reports the estimated post-1990 divergence in damages to GDP between treated and control countries using equation (A.1). The interaction term $Treated_i \times Post_t$ is positive, indicating that the shock has led to a substantial increase in damages to GDP in the treated group. The treated group suffer 0.119% of GDP more in damages per year on average over 1990-2000 than the control group. This is significant at the 10% significance level. In terms of magnitude, this is much greater than the average damages to GDP experienced by all countries in the 1990 to 2000, which was 0.08% of GDP per year.

This corroborates the other channel of the model, namely that countries which experience extreme weather events suffer from a constraint on their resources, proxied here by greater damages to GDP.

A.5 Robustness Checks

We conduct the following robustness checks that vary (i) the definition of an extreme episode (alternative cutoffs), (ii) the time horizon, (iii) the outcome definition (laws versus laws & policies; cumulative versus annual counts) (iv) placebo treatment years and (v) the treatment group to check that the results aren’t being driven by outliers. The qualitative finding, a relative post-episode slowdown in mitigation policy output, is broadly stable across these exercises.


The first robustness test in column 1 uses as the dependent variable the annual number of mitigation laws rather than the cumulative number of mitigation laws passed. The coefficient of the interaction term is still significant at the 5% level. After the 1990 shock, the treated countries have 0.068 fewer laws on average each year than the control group.

The second robustness test in column 2 changes the time frame to include a longer period going from 1980 to 2010, rather than our initial shorter period from 1980 to 2000¹³. As can be seen the coefficient on the interaction term remains significant at the 5% significant level.

The third robustness test changes the identification strategy. It identifies the extreme weather shock as years where there are more than 5σ above the country mean. Table A.8 lists these shocks by country and year. Using the 5σ cutoff decreases the number of countries that experience a shock to just Finland

¹³The significance of the interaction term when extending the horizon beyond 2010 may not necessarily imply that the impact of the 1990 shock continues to grow or re-emerge over time. Rather, it could reflect the persistence of the initial post-shock divergence between treated and control countries, which remains visible in cumulative terms even after two decades. Alternatively, the result might simply stem from the increased statistical power associated with a longer panel.

Figure A.2: Impact of climate disaster shock on annual damages to GDP



Note: 1990 is identified as the only year where the number of climate damages 4σ above the country mean. Treated group = Finland, Netherlands, Luxembourg and Denmark. Control group = Austria, Belgium, France, Germany Ireland, Italy, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom.

and Luxembourg. Figure A.3 below presents the plots of the treated and control group.

Table A.3: Robustness: The impact of 1990 disaster shock on mitigation laws

	(1)	(2)	(3)	(4)	(5)	(6)
main						
Interaction	-0.068** (0.032)	-0.893** (0.011)	-0.338*** (0.003)	0.205 (0.228)		
No climate disaster lag 1					-0.006 (0.949)	-0.018 (0.856)
No climate disaster lag 2					0.065 (0.508)	0.073 (0.464)
No climate disaster lag 3					0.098 (0.317)	0.093 (0.352)
No climate disaster lag 4					0.128 (0.186)	0.144 (0.156)
No climate disaster lag 5					-0.238** (0.022)	-0.266** (0.028)
N	336	496	336	315	576	576

P-value in parentheses | * 0.10 **0.05 ***0.01 | Country + Year fixed effects included


Columns 3 of table A.3 presents the results. As can be seen, the coefficient is -0.338, almost the same as when the 4σ cutoff is used, and is significant at the 1% significant level.

The fourth robustness test changes the identification strategy so that the cutoff is 3σ above the mean. Table A.9 lists these shocks by country and year.

Sweden is dropped from the sample as it experiences two shocks over the time period. All the other shocks also occur in 1990, however when defined in this way, the coefficient becomes insignificant. We interpret this results as suggesting that the shocks need to be sufficiently strong enough in order to have the impact on mitigation laws.

The fifth and sixth robustness test estimate a two-way fixed effects regression equation, given its use in the empirical literature Rowan, 2022. Given that we are now estimating the impact of multiple shocks over many years, we change the dependent variable to the frequency (rather than the cumulative frequency) of laws. This is different to the diff in diff, where there is one shock and we want to understand its total cumulative effect. Given that the dependent variable is a count variable with many potential zeros, we use two types of fixed effects regression, a Poisson fixed effects regression (column 4) and a zero-inflated negative binomial regression (column 5). We include 5 lags given the potential in the baseline for the shock to have an impact on laws into the future. Column 4 and 5 presents the results for laws and policies respectively. The fifth lag is significant and negative at 5% level for both estimations. All other lags are insignificant. Increasing the number of climate disasters a country experiences by 1,

Figure A.3: Robustness: Impact of climate disaster shock on cumulative no. mitigation laws with 5σ cutoff



Note: Shock is identified as years where the number of climate damages 5σ above the country mean. Full list of shocks is in table A.8 in the Appendix.

leads to a decline in the number of climate laws passed by -0.238/-0.266 after 5 years. The number of observations increases for this specification as we estimate this now over the whole period from 1980 to 2017, as we include all disasters and not just extreme shocks. While these effects are more muted, there is still some evidence that any shock can lead to a decline in the number of mitigation laws. This is to be expected given that not all these disasters will act as focusing events, particularly compared to the 1990 shock, and therefore the results are expected to be less significant.

The seventh set of robustness checks presented in Table A.4 estimates a placebo test by varying the start year of the treatment from 1990. Specifically, column 2 examines the effect of a placebo shock beginning in 1992 on the cumulative number of mitigation laws, column 3 considers the impact if the shock commenced in 1993, and so forth. The primary objective of this placebo test is to assess whether the observed impact of the 1990 shock genuinely drives the results. If the placebo tests (columns 2 to 7) yield statistically significant coefficients, it would imply that the 1990 shock may not be responsible for the observed differences between the treated and control groups. The dependent variable used in these tests is the cumulative number of laws and policies.

Table A.4: Placebo test: impact of shock on cumulative laws and policy with different shock start dates (1990 to 1996)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Start treatment year	1990	1991	1992	1993	1994	1995	1996
Interaction	-0.606* (0.089)	-0.619* (0.089)	-0.611* (0.099)	-0.615 (0.116)	-0.595 (0.152)	-0.586 (0.203)	-0.590 (0.269)
N	336	336	336	336	336	336	336

P-value in parentheses | * 0.10 **0.05 ***0.01 | Country and Year fixed effects are included

The coefficients become statistically insignificant after 1993, with the pattern of insignificance persisting through 1996. While the results suggest that 1990 is not the sole year exhibiting significance, they also indicate that the further the hypothetical shock deviates from 1990, the less significant the coefficients become. This finding underscores that the 1990 shock likely captures a genuine effect. To summarise, we find that unexpected climate disasters significantly reduce the number of climate mitigation laws implemented. This result is robust across several identification strategies and controls.

The last set of robustness tests checks whether the results are being driven by a single country, particularly as the number of treated countries is four. To check this, we estimate the baseline regression, removing one country at a time in Table A.5. The results remain significant in all specifications.

Table A.5: Robustness: Leave-one-out test excluding one treated country at a time

	(1)	(2)	(3)	(4)
	Excl. NLD	Excl. FIN	Excl. DNK	Excl. LUX
ATET				
interaction	-0.364*** (0.002)	-0.252* (0.069)	-0.252* (0.069)	-0.252* (0.069)
N	336	336	336	336

P-value in parentheses | * 0.10 **0.05 ***0.01 | Country and Year fixed effects are included

A.6 The 1990 Windstorm Episode

Table A.6: Major European Windstorms in 1990

Storm Name	Dates	Main Countries Affected	Notable Impacts
Burns' Day Storm (Daria)	25–26 Jan 1990	UK, Netherlands, Denmark, Luxembourg, Germany	Winds up to 230 km/h; 97 fatalities across Europe; severe damage to infrastructure and forests; widespread power outages.
Herta	1–6 Feb 1990	UK, France, Germany, Belgium, Netherlands	Significant wind damage; insured losses estimated at \$1.5 billion (2012 USD); disruption to transport and utilities.
Judith	7–8 Feb 1990	Western Europe	Moderate storm with localized damage; limited widespread impact.
Nana	11–12 Feb 1990	Western Europe	Part of the 1990 storm sequence; caused minor structural and tree damage.
Otilie	13–14 Feb 1990	Western Europe	Minor to moderate impacts; damage to roofs, trees, and power lines.
Polly	14–15 Feb 1990	Western Europe	Localized wind damage; part of prolonged storm activity in February.
Vivian	25–28 Feb 1990	Netherlands, Finland, Germany, Switzerland	Winds up to 268 km/h; 64 fatalities; extensive forest damage (esp. in Germany and Switzerland); transport disruptions.
Wiebke	28 Feb – 1 Mar 1990	Germany, Netherlands, France, Switzerland	Major forest destruction; severe economic damage; insured losses of \$1.4 billion (2012 USD); followed closely after Vivian.

Sources: EM-DAT – The International Disaster Database; European Windstorm Database; ESWD – European Severe Weather Database.

A.7 Extreme-Event Definitions

Table A.7: Periods where frequency of disasters is more than 4σ above country mean

Country	Date	No. Disaster	Mean No. Disaster	σ No. Disaster
Denmark	1990	3	.372093	.6554989
Finland	1990	2	.0697674	.337734
Luxembourg	1990	6	.3023256	.9644856
Netherlands	1990	6	.8604651	1.125069

Table A.8: Periods where frequency of disasters is more than 5σ above country mean

Country	Date	No. Disaster	Mean No. Disaster	σ No. Disaster
Finland	1990	2	.0697674	.337734
Luxembourg	1990	6	.3023256	.9644856


Table A.9: Periods where frequency of disasters is more than 3σ above country mean

Country	Date	No. Disaster	Mean No. Disaster	σ No. Disaster
Sweden	1990	2	.255814	.5386502
Norway	1990	2	.3023256	.5133867
Austria	1990	5	1.116279	1.13828
Belgium	1990	6	1.255814	1.39886
Ireland	1990	3	.5348837	.6305265
Netherlands	1990	6	.8604651	1.125069
Denmark	1990	3	.372093	.6554989
Finland	1990	2	.0697674	.337734
Luxembourg	1990	6	.3023256	.9644856
Sweden	2005	2	.255814	.5386502
Spain	2019	8	2.069767	1.486372

A.8 Summary Statistics and Ancillary Evidence


	Obs	Mean	σ	Min	Max
Annual frequency mitigation laws Climate Laws Dataset	656	0.143	0.414	0.000	3.000
Annual climate damages to GDP % EM-DAT	656	0.056	0.207	0.000	2.895
Annual frequency of climate related disasters EM-DAT	656	1.220	1.629	0.000	10.000
Coal gas and oil rents % GDP World Bank	654	0.613	1.800	0.000	12.248
Number of people (1000s) World Bank	656	23933.321	25870.970	364.150	83160.871
Left-right index of chief minister (1=right, 2=centre, 3=left)	568	1.879	0.932	0.000	3.000
Democracy index (0–1) Uni Gothenberg Var of Demo	656	0.877	0.024	0.797	0.924
GDP per capita constant 2015 USD (1000s) World Bank	656	41.287	19.803	10.744	112.418
Annual GHG in CO2 equiv (million tonnes) Our World in Data	656	244.979	283.422	8.651	1299.919
Share of people who believe GW is a serious concern EM-DAT	71	0.591	0.151	0.296	0.866

Figure A.4: Damages to GDP from Extreme Weather Events: Treatment and Control Group (% of GDP)




Note: 1990 is identified as the only year where the number of climate disasters exceeds 4σ above the country mean. Treated group = Finland, Netherlands, Luxembourg and Denmark. Control group = Austria, Belgium, France, Germany, Ireland, Italy, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom.

Figure A.5: Average number of climate related disasters 1980-2020



Note: Data from EMDAT database. Climate related disasters are defined as drought, extreme temperature, flood, landslide, storm and wildfire.

Figure A.6: Average number of mitigation laws per year 1980-2020



Note: Data from Climate Laws of the World. Mitigation laws are all laws that only refer to mitigation.

A.9 Eurobarometer Data Construction

It is not possible to conduct the above exercise using data on attitudes to global warming due to a lack of consistent data over this period. We do however construct a climate concern variable from Eurobarometer surveys conducted in 1989 and 1991. Respondents are classified as “seriously concerned” about global warming if they rank it among their top environmental concerns. The pre/post comparison uses the survey waves immediately flanking the 1990 shock to minimise confounding from other time trends.

Figure A.7a, using data from Eurobarometer¹⁴, shows that the share of people seriously concerned about global warming increased in all countries after the shock, and increased most in countries with the highest damages to GDP. A line of best fit shows a positive and significant (at the 5% level) relationship between damages and heightened concern over global warming, although this is being driven by Luxembourg.¹⁵ In Luxembourg the share of people seriously concerned about global warming increased on average by 30 percentage points between 1989 and 1991, while the 1990 extreme weather shock caused 3% of climate damages to Luxembourg’s GDP.

Yet this increase in concern was not associated with a corresponding increase in mitigation laws in the most affected countries. Figure A.7b, using data from Climate Laws of the World, shows that the number of mitigation laws, while increasing across all countries, increased least in those with the highest damages to GDP. Luxembourg, which experienced the largest shock, was the only country that did not increase its mitigation laws over this period.

These motivating patterns suggest an organising framework with two opposing forces. Disasters may increase perceived climate risk and public demand for action (the signal channel), but at the same time they can tighten contemporaneous constraints and shift governmental priorities toward reconstruction and crisis governance (the output channel). To document the second leg of this trade-off more systematically, we now turn from descriptive scatter plots to a difference-in-differences exercise based on the same 1990 episode.

B Proofs

B.1 Preliminaries


Throughout, write $u_n \equiv u(AK_1)$, $u_d \equiv u((1 - \delta)AK_1)$, $|\Delta u| \equiv u_n - u_d > 0$, and

$$\begin{aligned} \Delta u &\equiv u_d - u_n < 0, \\ \Delta u_K &\equiv (1 - \delta)Au'((1 - \delta)AK_1) - Au'(AK_1) = Au'(AK_1)[(1 - \delta)^{1-\sigma} - 1] \geq 0 \text{ for } \sigma \geq 1, \end{aligned}$$

¹⁴See Appendix A.9 for a detailed construction of this variable.

¹⁵The difference-in-difference regressions below, however, are significant even when Luxembourg is not part of the sample.

Figure A.7: Changes in concern about global warming and mitigation laws vs damages to GDP from the 1990 shock



Notes: In both panels, the x-axis shows $\log(1 + \text{damages to GDP from climate-related extreme weather events in 1990})$. We use logs given the potentially thick-tailed nature of climate shocks to GDP. Data are from EM-DAT. Panel (a) plots the change in the average share of people in each country who are seriously concerned about global warming from just before the shock (1989) to just after (1991), using Eurobarometer data. Panel (b) plots the change in the average annual frequency of mitigation laws implemented between the post-shock period (1991–2010) minus the pre-shock period (1980–1989), using data from Climate Laws of the World. Data are available only for a select number of countries.

so that $\partial|\Delta u|/\partial K_1 = -\Delta u_K \leq 0$. Define $g \equiv e^{\Theta|\Delta u|} - 1 > 0$ and $Z \equiv 1 + \hat{p}_1 g \geq 1$, so that the objects of (8)–(10) are $\Phi = \Theta^{-1} \ln Z$, $\mathcal{M} = g/(\Theta Z)$, and $\pi = \hat{p}_1 e^{\Theta|\Delta u|}/Z$, and the objective under the exact certainty equivalent is

$$V_0 = u(C_0) + \beta [u_n - \Phi(\hat{p}_1, |\Delta u|; \Theta)].$$

Five identities are used repeatedly; all follow from direct differentiation together with $Z - \hat{p}_1 g = 1$ and $Z - \hat{p}_1 e^{\Theta|\Delta u|} = 1 - \hat{p}_1$:

$$\begin{aligned} \frac{\partial \Phi}{\partial \hat{p}_1} &= \mathcal{M}, & \frac{\partial \mathcal{M}}{\partial \hat{p}_1} &= -\frac{g^2}{\Theta Z^2} = -\Theta \mathcal{M}^2, & \frac{\partial \Phi}{\partial |\Delta u|} &= \pi, \\ \frac{\partial \pi}{\partial |\Delta u|} &= \Theta \pi(1 - \pi), & \frac{\partial \mathcal{M}}{\partial |\Delta u|} &= \frac{\partial \pi}{\partial \hat{p}_1} = \frac{e^{\Theta|\Delta u|}}{Z^2} = \frac{\pi(1 - \pi)}{\hat{p}_1(1 - \hat{p}_1)} \equiv \Pi > 0. \end{aligned}$$

As $\Theta \rightarrow 0$: $\Phi \rightarrow \hat{p}_1 |\Delta u|$, $\mathcal{M} \rightarrow |\Delta u|$, $\pi \rightarrow \hat{p}_1$, $\Pi \rightarrow 1$, recovering expected utility. To first order in Θ , $\mathcal{M} = |\Delta u| [1 + \frac{\Theta}{2}(1 - 2\hat{p}_1)|\Delta u|] + O(\Theta^2)$ and $\Pi = 1 + \Theta(1 - 2\hat{p}_1)|\Delta u| + O(\Theta^2)$, which deliver the mean–variance forms quoted in the text. The leading-order expansions in δ , used where indicated, are $|\Delta u| = q\delta + \frac{\sigma q}{2}\delta^2 + O(\delta^3)$ with $q \equiv AK_1 u'(AK_1)$, and $\Delta u_K = (\sigma - 1)\delta Au'(AK_1) + O(\delta^2)$.

B.2 First-Order Conditions

Mitigation affects C_0 at rate -1 and $\hat{p}_1$ at rate $-\alpha\kappa_1$ through $\hat{H}_1$; reconstruction affects C_0 at rate -1 and K_1 at rate $+1$. Using $\partial\Phi/\partial\hat{p}_1 = \mathcal{M}$ and

$$\frac{\partial}{\partial K_1} [u_n - \Phi] = Au'(AK_1) + \pi \Delta u_K = (1 - \pi) Au'(AK_1) + \pi(1 - \delta)Au'((1 - \delta)AK_1) = \mathcal{B},$$

the first-order conditions are

$$\frac{\partial V_0}{\partial m_0} = -u'(C_0) + \beta\alpha\kappa_1 \mathcal{M} = 0, \tag{FOC-m}$$

$$\frac{\partial V_0}{\partial r_0} = -u'(C_0) + \beta \mathcal{B} = 0, \tag{FOC-r}$$

with $\mathcal{B}$ the tilted-mean marginal product of capital defined in Assumption 2.

B.3 Continuous Embedding

For the comparative statics in Theorem 2 and Proposition 2, the binary realisation is embedded in the family of problems $\{\mathcal{P}(s)\}_{s \in [0,1]}$ obtained by setting $D_0 = s$ in (6) and (5), as in (15). The feasible set at exposure s is

$$\mathcal{F}(s) = \{(m_0, r_0) \in \mathbb{R}_+^2 : m_0 + r_0 < AK_0(1 - \delta s)\},$$

and the shock derivatives are $\partial C_0/\partial s = -\delta AK_0$ and $\partial \hat{H}_1/\partial s = (1+\theta)\psi_H$, both constant since $\mathbb{E}[D_0] = p(H_0)$ is predetermined. The maintained regularity conditions and Assumptions 1–2 hold at every $s \in [0, 1]$, so Theorem 1 applies verbatim at each s . Because V_0 is C^2 in (m_0, r_0, s) and $\Delta_H > 0$, the implicit function theorem gives a continuously differentiable policy path with

$$\begin{pmatrix} dm_0^*/ds \\ dr_0^*/ds \end{pmatrix} = -H(s)^{-1} \begin{pmatrix} V_{ms} \\ V_{rs} \end{pmatrix}, \quad H(s) = \begin{pmatrix} V_{mm} & V_{mr} \\ V_{mr} & V_{rr} \end{pmatrix},$$

all evaluated at $(m_0^*(s), r_0^*(s); s)$. The fundamental theorem of calculus then delivers (16).

B.4 Second Derivatives and Hessian

Differentiating (FOC-m)–(FOC-r) and using the identities above,

$$\begin{aligned} V_{mm} &= u''(C_0) - \beta\alpha^2[\kappa_2\mathcal{M} - \Theta\kappa_1^2\mathcal{M}^2], \\ V_{mr} &= u''(C_0) - \beta\alpha\kappa_1\Pi\Delta u_K, \\ V_{rr} &= u''(C_0) + \beta[(1-\pi)A^2u''(AK_1) + \pi(1-\delta)^2A^2u''((1-\delta)AK_1) - \Theta\pi(1-\pi)(\Delta u_K)^2]. \end{aligned}$$

Three sign facts hold exactly; write $c \equiv \sigma C_0^{-\sigma-1} = |u''(C_0)|$ and recall $\mathcal{G}$ and $\mathcal{W}$ from Assumption 1. First, $V_{rr} = -(c + \mathcal{G}) < 0$ unconditionally for $\Theta \geq 0$. Second, $V_{mr} = -(c + \mathcal{W}) \leq u''(C_0) < 0$ for $\sigma \geq 1$, with equality in the log limit, where $\Delta u_K = 0$. Third, Assumption 1 delivers strict concavity exactly: since the mean-channel curvature only adds concavity, $-V_{mm} = c + \beta\alpha^2\kappa_2\mathcal{M} - \beta\Theta\alpha^2\kappa_1^2\mathcal{M}^2 \geq c - \beta\Theta\alpha^2\kappa_1^2\mathcal{M}^2$, which is strictly positive because the product in (11) exceeds $(c + \mathcal{W})^2 \geq c^2 > 0$ while $c + \mathcal{G} > 0$; and

$$\Delta_H = (-V_{mm})(-V_{rr}) - V_{mr}^2 \geq (c - \beta\Theta\alpha^2\kappa_1^2\mathcal{M}^2)(c + \mathcal{G}) - (c + \mathcal{W})^2 > 0$$

by (11), with no approximation. Assumption 1 also forces $\mathcal{G} > \mathcal{W}$: were $\mathcal{W} \geq \mathcal{G}$, then $(c + \mathcal{W})^2 \geq (c + \mathcal{G})^2$ and (11) would require $c - \beta\Theta\alpha^2\kappa_1^2\mathcal{M}^2 > c + \mathcal{G}$, which is impossible. Hence $V_{rr} - V_{mr} = -(\mathcal{G} - \mathcal{W}) < 0$ exactly. Combining (11) with the first-order conditions at an interior optimum yields the curvature bound (13) stated in the text: (FOC-m) gives $\beta\alpha\kappa_1 = u'(C_0)/\mathcal{M}$, so $\mathcal{W} = u'(C_0)\Pi\Delta u_K/\mathcal{M} = \Pi(\sigma - 1)u'(C_0)/K_1 + O(\delta)$ by the leading-order expansions above, while (FOC-r) gives $u'(C_0) = \beta\mathcal{B} = \beta Au'(AK_1) + O(\delta)$, so $\mathcal{G} = \sigma C_0^{-\sigma}/K_1 + O(\delta + \Theta\delta^2)$. Hence $\mathcal{W}/c = \Pi(\sigma - 1)C_0/(\sigma K_1) + O(\delta)$ and $\mathcal{G}/c = C_0/K_1 + O(\delta + \Theta\delta^2)$; substituting into (11), dropping the $O(\Theta)$ saturation term and letting $\Pi \rightarrow 1$, the condition reduces to $2(\sigma - 1)/\sigma - 1 + (\sigma - 1)^2(C_0/K_1)/\sigma^2 < 0$, that is, $C_0/K_1 < \sigma(2 - \sigma)/(\sigma - 1)^2$, which is (13). Two further exact consequences at an interior optimum are worth recording. First, (FOC-m) gives $\beta\Theta\alpha^2\kappa_1^2\mathcal{M}^2 = \Theta u'(C_0)^2/\beta$ with no approximation. Because the right-hand side of (11) satisfies $(c + \mathcal{W})^2 \geq c^2 > 0$ while $c + \mathcal{G} > 0$, the first factor of (11) must itself be strictly positive, $c > \beta\Theta\alpha^2\kappa_1^2\mathcal{M}^2$; substituting the interior-optimum identity and $u'(C_0) = C_0^{-\sigma}$ yields $\Theta < \sigma\beta C_0^{\sigma-1}$: risk sensitivity faces an effective ceiling in the units of u , and $\Theta > 0$ tightens (11) relative to the

expected-utility case. Concavity of the mitigation margin alone, $-V_{mm} > 0$, is equivalent at an interior optimum to the weaker bound $\Theta < \sigma\beta C_0^{\sigma-1} + \beta^2\alpha^2\kappa_2\mathcal{M}C_0^{2\sigma}$; the mean-channel slack in that bound is what Assumption 1 does not use. Second, the same substitutions pin the cross-wedge share used in the severity analysis by curvature alone at leading order: $\tilde{w} \equiv \mathcal{W}/(\mathcal{G}-\mathcal{W}) = \Pi(\sigma-1)/(\sigma-\Pi(\sigma-1)) + O(\delta)$, independent of C_0/K_1 , and $\tilde{w} = 0$ in the log limit.

B.5 Proof of Proposition 1

Holding m_0 fixed in (5), direct differentiation gives $\partial\hat{H}_1/\partial D_0 = (1+\theta)\psi_H$. Since $\hat{p}_1 = p(\hat{H}_1)$ and $p'(\hat{H}_1) > 0$ by the maintained monotonicity of p , the chain rule gives $\partial\hat{p}_1/\partial D_0 = p'(\hat{H}_1)(1+\theta)\psi_H > 0$ whenever $\psi_H > 0$. At $\theta = 0$ the response is the rational conditional response ψ_H ; for $\theta > 0$ it is strictly larger. $\square$

B.6 Proof of Theorem 1

Write $\bar{Y} \equiv AK_0(1-\delta D_0)$, so that $\mathcal{F} = \{(m_0, r_0) \in \mathbb{R}_+^2 : m_0 + r_0 < \bar{Y}\}$ and $C_0 = \bar{Y} - m_0 - r_0 > 0$ on $\mathcal{F}$.

Existence. The set $\mathcal{F}$ is not closed, so existence requires a compactness argument. On the closure of $\mathcal{F}$ the continuation value is bounded above: $K_1 = (1-d)K_0 + r_0 \in [(1-d)K_0, (1-d)K_0 + \bar{Y}]$, so $u(AK_1) \leq u(A[(1-d)K_0 + \bar{Y}])$, and $\Phi \geq 0$ since $Z \geq 1$; hence $\beta[u(AK_1) - \Phi] \leq B$ for a finite constant B . Under CRRA felicity with $\sigma \geq 1$, $u(C_0) \rightarrow -\infty$ as $C_0 \rightarrow 0^+$, so $V_0 \leq u(C_0) + B \rightarrow -\infty$ uniformly as $m_0 + r_0 \rightarrow \bar{Y}$. The reference point $(0, 0) \in \mathcal{F}$ gives a finite value, so there is $\varepsilon > 0$ such that $V_0 < V_0(0, 0)$ whenever $C_0 < \varepsilon$, and the supremum of V_0 over $\mathcal{F}$ equals its supremum over the compact set $\mathcal{F}_\varepsilon \equiv \{(m_0, r_0) \in \mathbb{R}_+^2 : m_0 + r_0 \leq \bar{Y} - \varepsilon\}$, on which V_0 is continuous. A maximiser exists by Weierstrass.

Uniqueness. By the second-derivative facts above, Assumption 1 delivers $V_{mm} < 0$ and $\Delta_H > 0$ at every point of $\mathcal{F}$, so the Hessian of V_0 is negative definite everywhere and V_0 is strictly concave on the convex set $\mathcal{F}$. A strictly concave function on a convex set has at most one maximiser.

Interiority. Let (m_0^*, r_0^*) be the maximiser and suppose $m_0^* = 0$. The direction $(1, 0)$ is feasible at $(0, r_0^*)$, since the upper constraint is slack there, and the one-sided directional derivative in that direction is the left-hand side of (FOC-m), $\partial V_0/\partial m_0 = -u'(C_0) + \beta\alpha\kappa_1\mathcal{M}$, which is strictly positive at the candidate maximiser by Assumption 2, condition (12), since the maximiser lies in the relevant upper-contour set. This contradicts optimality: equivalently, it violates the Karush–Kuhn–Tucker requirement $\partial V_0/\partial m_0 \leq 0$ at a maximiser whose non-negativity constraint binds. The same argument with the direction $(0, 1)$, the left-hand side of (FOC-r), and the second part of Assumption 2 rules out $r_0^* = 0$. Hence $m_0^*, r_0^* > 0$, the maximiser is interior, the first-order conditions (FOC-m)–(FOC-r) hold there, and by strict concavity this interior critical point is the unique global maximum. The argument applies verbatim at every exposure $s \in [0, 1]$ of the embedding, with $\bar{Y} = AK_0(1 - \delta s)$. $\square$

B.7 Proof of Theorem 2

Fix $s \in [0, 1]$. Differentiating the first-order system with respect to s and applying Cramer's rule:

$$\frac{dm_0^*(s)}{ds} = \frac{-V_{ms}V_{rr} + V_{rs}V_{mr}}{\Delta_H}.$$

The shock enters each condition through $C_0(s)$, with $\partial C_0/\partial s = -\delta AK_0$, and through $\hat{H}_1(s)$, with $\partial \hat{H}_1/\partial s = (1 + \theta)\psi_H$. For the mitigation condition (FOC-m), the consumption channel contributes $u''(C_0)\delta AK_0 = -O$; for the belief channel, $|\Delta u|$ depends only on (K_1, δ) , so only κ_1 and $\hat{p}_1$ vary with $\hat{H}_1$, and

$$\frac{\partial}{\partial \hat{H}_1} [\beta \alpha \kappa_1 \mathcal{M}] = \beta \alpha \left[\kappa_2 \mathcal{M} + \kappa_1^2 \frac{\partial \mathcal{M}}{\partial \hat{p}_1} \right] = \beta \alpha [\kappa_2 \mathcal{M} - \Theta \kappa_1^2 \mathcal{M}^2];$$

multiplying by $(1 + \theta)\psi_H$ gives exactly S in (18), so $V_{ms} = S - O$. For the reconstruction condition (FOC-r), the consumption channel again contributes $-O$; the belief channel runs through $\hat{p}_1$ alone, with

$$\frac{\partial \mathcal{B}}{\partial \hat{p}_1} = \frac{\partial \pi}{\partial \hat{p}_1} [(1 - \delta)Au'((1 - \delta)AK_1) - Au'(AK_1)] = \Pi \Delta u_K,$$

so the belief channel contributes $(1 + \theta)\psi_H \beta \kappa_1 \Pi \Delta u_K = S_r$ and $V_{rs} = S_r - O$. Substituting:

$$\frac{dm_0^*(s)}{ds} = \frac{-(S - O)V_{rr} + (S_r - O)V_{mr}}{\Delta_H} = \frac{1}{\Delta_H} [-V_{rr}S + V_{rr}O - V_{mr}O + V_{mr}S_r],$$

which is (17). Since s was arbitrary, the decomposition holds at every point of the path. $\square$

B.8 Discrete-Effect Statement

The IFT delivers a C^1 path, so (16) holds by the fundamental theorem of calculus. The integrand equals Δ_H^{-1} times the bracketed expression in (17), with $\Delta_H > 0$. If the bracket is strictly negative at every s , the integrand is a continuous strictly negative function on $[0, 1]$, and its integral is strictly negative. The positive case is symmetric. $\square$

B.9 Proof of Proposition 2

(i) Fix s . By Theorem 2 and $\Delta_H > 0$, $dm_0^*(s)/ds < 0$ iff

$$V_{rr}(O - S) + V_{mr}(S_r - O) < 0.$$

Dividing by $V_{rr} < 0$ reverses the inequality and rearranging gives $\tilde{\Lambda}(s) > S(s)$. (ii) If the inequality holds at every s , part (i) and the discrete-effect statement following Theorem 2 gives $\Delta m_0^* < 0$. $\square$

For the comparative statics in the list following Proposition 2, note that $V_{rr} - V_{mr} = -(\mathcal{G} - \mathcal{W}) < 0$ exactly under Assumption 1, by the sign facts above, so $\Lambda = O(V_{rr} - V_{mr})/V_{rr} > 0$ as the ratio of

two negative quantities; at leading order in δ , $V_{rr} - V_{mr} = \beta A^2 u''(AK_1) + O(\delta)$, the form used in the leading-order readings.

B.10 Derivations for the Severity Discussion

All objects are evaluated along the embedding at the disaster endpoint $s = 1$, with the local values $(\hat{p}_1, \kappa_1, \kappa_2)$ of the probability function treated as parameters of the evaluation point. The elasticities are direct-effect elasticities, computed holding the policy pair fixed at its optimum while δ varies. The re-optimisation terms, $(\partial \ln f / \partial m_0) dm_0^* / d\delta + (\partial \ln f / \partial r_0) dr_0^* / d\delta$ for $f \in \{S, \tilde{\Lambda}\}$, enter the total elasticities at the same $O(\delta)$ order and involve third derivatives of p and u that the maintained assumptions do not sign; the elasticities are therefore reported as direct severity effects. At $\chi = 0$, where the direct margin is itself $O(\delta)$, the unsigned feedback terms can reverse the ranking, so the comparison there is a statement about the severity channel in isolation. With $|\Delta u| = q\delta + \frac{\sigma q}{2}\delta^2 + O(\delta^3)$, $q \equiv AK_1 u'(AK_1)$, and $\mathcal{M} = |\Delta u| + \Theta(\frac{1}{2} - \hat{p}_1)|\Delta u|^2 + O(|\Delta u|^3)$, the diagnostic bracket in (18) expands as

$$\kappa_2 \mathcal{M} - \Theta \kappa_1^2 \mathcal{M}^2 = a\delta - \tilde{b}\delta^2 + O(\delta^3), \quad a = \kappa_2 q, \quad \tilde{b} = \Theta \kappa_1^2 q^2 - \frac{\kappa_2}{2} [\sigma q + \Theta q^2 (1 - 2\hat{p}_1)],$$

so that $S = (1 + \theta) \bar{\psi}_H \delta^\chi \beta \alpha (a\delta - \tilde{b}\delta^2 + O(\delta^3))$ and

$$\varepsilon_S = 1 + \chi - \frac{(\tilde{b}/a)\delta}{1 - (\tilde{b}/a)\delta}.$$

The signal saturates at second order, $\tilde{b} > 0$, iff $\Theta q \kappa_1^2 > \frac{\kappa_2}{2} [\sigma + \Theta q (1 - 2\hat{p}_1)]$: prudence convexifies the mean channel through the σq term, so saturation is a curvature condition rather than an unconditional fact, satisfied when κ_2 is moderate. For the resource side, write $c \equiv \sigma C_0^{-\sigma-1}$, so that $V_{mr} = -(c + \mathcal{W})$, $V_{rr} = -(c + \mathcal{G})$, and $\Lambda = O(\mathcal{G} - \mathcal{W}) / (c + \mathcal{G})$. Under the direct-effect convention the severity perturbation moves C_0 (hence c and O), $|\Delta u|$, Δu_K , π , and Π , holding the primitives and the policy pair fixed. Then $O = c \delta AK_0$ with $\partial C_0 / \partial \delta = -AK_0$ at $s = 1$ gives $\varepsilon_O = 1 + (\sigma + 1) \delta AK_0 / C_0$ and $\varepsilon_c = (\sigma + 1) \delta AK_0 / C_0$; $\mathcal{G} = \beta A^2 |u''(AK_1)| + O(\delta)$ is severity-flat at leading order, $\varepsilon_{\mathcal{G}} = O(\delta)$; and $\mathcal{W} = \beta \alpha \kappa_1 \Pi \Delta u_K$ inherits $\varepsilon_{\mathcal{W}} = 1 + O(\delta)$ from $\Delta u_K = (\sigma - 1) \delta A u'(AK_1) + O(\delta^2)$. Hence

$$\varepsilon_\Lambda = \varepsilon_O + \frac{\varepsilon_{\mathcal{G}} \mathcal{G} - \varepsilon_{\mathcal{W}} \mathcal{W}}{\mathcal{G} - \mathcal{W}} - \frac{\varepsilon_c c + \varepsilon_{\mathcal{G}} \mathcal{G}}{c + \mathcal{G}} = 1 + \omega(\sigma + 1) \frac{\delta AK_0}{C_0} - \tilde{w} (1 + O(\delta)), \quad \tilde{w} \equiv \frac{\mathcal{W}}{\mathcal{G} - \mathcal{W}},$$

with $\omega \equiv \beta A^2 u''(AK_1) / V_{rr} = \mathcal{G} / (c + \mathcal{G}) + O(\delta) \in (0, 1)$. The wedge term is not of order δ : by the interior-optimum substitutions in the second-derivative subsection above, $\tilde{w} = \Pi(\sigma - 1) / (\sigma - \Pi(\sigma - 1)) + O(\delta)$, which is of order $\sigma - 1$, independent of C_0 / K_1 at leading order, and vanishes identically in the log-felicity limit, where $\mathcal{W} \equiv 0$ and $\varepsilon_\Lambda = 1 + \omega(\sigma + 1) \delta AK_0 / C_0$ at leading order. Write $\tilde{\Lambda} = \Lambda + T$ with $T \equiv (V_{mr} / V_{rr}) S_r \geq 0$, so that $\varepsilon_{\tilde{\Lambda}} = \varsigma \varepsilon_\Lambda + (1 - \varsigma) \varepsilon_T$ with $\varsigma \equiv \Lambda / \tilde{\Lambda} \in (0, 1]$. Since $T \propto \psi_H(\delta) \kappa_1 \Pi \Delta u_K$ scales as $\delta^{1+\chi}$ at leading order, $\varepsilon_T = 1 + \chi + O(\delta)$ with the $O(\delta)$ correction free of χ , and $T / \Lambda = O(\delta^\chi)$ with leading coefficient proportional to $(1 + \theta) \bar{\psi}_H \kappa_1 \Pi(\sigma - 1)$, so the reconstruction share $1 - \varsigma$ is small

at $\chi = 0$ precisely when $\sigma = 1$, the transmission $\bar{\psi}_H$, or the slope κ_1 is small. Under $\tilde{b} \geq 0$ and $\chi + \tilde{w} < \omega(\sigma + 1)\delta AK_0/C_0$,

$$\varepsilon_{\tilde{\Lambda}} - \varepsilon_S = \varsigma(\varepsilon_{\Lambda} - \varepsilon_S) + (1 - \varsigma)(\varepsilon_T - \varepsilon_S) = \varsigma \left[\omega(\sigma + 1) \frac{\delta AK_0}{C_0} - \tilde{w} - \chi + \frac{(\tilde{b}/a)\delta}{1 - (\tilde{b}/a)\delta} \right] + O((\sigma - 1)\delta + \delta^{1+\chi} + \delta^2),$$

where the bracket is strictly positive under the displayed condition and the residual collects the $O(\delta)$ parts of $\varepsilon_{\mathcal{W}}$ and $\varepsilon_{\mathcal{G}}$ together with $(1 - \varsigma)(\varepsilon_T - \varepsilon_S)$. The leading such term is the disaster-branch curvature drift of $\mathcal{G}$: since $\beta\pi(1 - \delta)^2 A^2 |u''((1 - \delta)AK_1)| = \beta\pi\sigma A^2 (AK_1)^{-\sigma-1} (1 - \delta)^{1-\sigma}$, direct differentiation gives $\varepsilon_{\mathcal{G}} = \pi(\sigma - 1)\delta + O(\delta^2)$, which enters ε_{Λ} with the order-one weight $1 + \tilde{w} - \omega$ and is therefore of the same order in δ as the margin, with a coefficient of order $\sigma - 1$; it is dominated precisely in the near-log regime $\sigma - 1 = O(\delta)$, where it is $O(\delta^2)$. Outside that regime these terms perturb the margin at its own order, which is why the comparison is a leading-order statement. Because $\tilde{w} = O(\sigma - 1)$ while the margin is $O(\delta)$, the condition has content only when $\sigma - 1$ is comparable to δ : the comparison is exact at leading order in the log-felicity limit, where $\tilde{w} = 0$, $T \equiv 0$, and $\varsigma = 1$; sufficient near it; and quantitative rather than sign-determined for σ appreciably above one. $\square$

C Additional Theoretical Derivations

C.1 Expected-Utility Benchmark

Set $\Theta = 0$ in (18). Since $\mathcal{M} \rightarrow |\Delta u|$, the saturation term disappears and $S^{EU} = (1 + \theta)\psi_H \beta \alpha \kappa_2 |\Delta u|$. The sign statements follow from $\psi_H \geq 0$, $\beta > 0$, $\alpha > 0$, $|\Delta u| > 0$, and $\kappa_2 = p''(\hat{H}_1) \geq 0$. $\square$

C.2 Diagnostic-Intensity Comparative Static

From (18), $S = (1 + \theta)\psi_H \beta \alpha B_S$. Holding the other local objects fixed, $\partial S / \partial \theta = \psi_H \beta \alpha B_S$. The derivative is positive if $B_S > 0$ and negative if $B_S < 0$. Since defocusing requires $\tilde{\Lambda} > S$ in Proposition 2, an increase in S shrinks the defocusing region and a decrease in S expands it. $\square$

C.3 Institutional-Design Implication

The defocusing condition is $\tilde{\Lambda} > S$, with $\tilde{\Lambda} = \Lambda + (V_{mr}/V_{rr})S_r$ and $\Lambda = O(V_{rr} - V_{mr})/V_{rr}$. Any change that lowers O lowers Λ holding the other local objects fixed; any change that lowers S_r lowers $\tilde{\Lambda}$ because $V_{mr}/V_{rr} > 0$; and any change that raises a positive S makes $\tilde{\Lambda} > S$ harder to satisfy. These changes shrink the set of primitives for which defocusing occurs. $\square$

D Dynamic Model

D.1 Diagnostic Expectations in the Dynamic Model

The two-period belief equation (5) specifies the diagnostic perception for a single realisation. The infinite-horizon model requires an operator, together with a timing convention for the impulse responses. Information and timing are as follows: at the start of period t the pair (K_t, H_t) and all past realisations are known, the disaster indicator D_t realises, policies (m_t, r_t) are chosen, and the Gaussian innovation η_t realises into H_{t+1} through (25).

Diagnostic beliefs operate on the voter's estimate of the hazard state. The *perceived hazard state* $\hat{H}_t$ follows

$$\hat{H}_{t+1} = \rho \hat{H}_t - \alpha m_t + \psi_H D_t + \eta_t + \theta[\psi_H(D_t - \bar{p}_t) + \eta_t], \quad \hat{H}_0 = H_0, \quad (\text{D.1})$$

the rational update of the perceived hazard state plus θ times the period- t hazard news, where $\bar{p}_t$ is the comparison probability of the diagnostic distortion. Under the distant-memory convention adopted here, $\bar{p}_t$ is the pre-disturbance norm: $\bar{p}_0 = p(H_0)$ at the impact date, so that (D.1) nests the two-period equation (5) at $t = 0$, and $\bar{p}_t = p(H^*)$ at the steady state. Because $\bar{p}_t$ is predetermined, $\partial \hat{H}_{t+1} / \partial \hat{H}_t = \rho$: the perceived hazard state inherits the persistence of the true hazard state. The operator $\mathbb{E}_t^{DE}$ in the Bellman equation is expectation under the diagnostic conditional law implied by (D.1): conditional on period- t information, $H_{t+1} \sim N(\rho \hat{H}_t - \alpha m_t + \psi_H D_t + \theta \psi_H(D_t - \bar{p}_t), \sigma_\eta^2)$ and $D_{t+1} \sim \text{Bernoulli}(p(\hat{H}_{t+1}))$. The voter's state is therefore $(K_t, \hat{H}_t)$; since $\hat{H}_t = H_t$ along any path without news, we keep the notation $V(K_t, H_t)$, with H_t understood as the perceived hazard state.

Both pieces of (D.1) follow from the representativeness distortion of Bordalo, N. Gennaioli, and Shleifer (2018) applied to the two components of the innovation. For the Gaussian component, the diagnostic density $f^\theta \propto f(\cdot | \mathcal{I}_t) [f(\cdot | \mathcal{I}_t) / f(\cdot | \mathcal{I}_t^-)]^\theta$, where the comparison conditional $\mathcal{I}_t^-$ carries the no-news mean and the same variance, is again Gaussian with variance σ_η^2 and mean shifted by θ times the forecast revision; this is the $\theta \eta_t$ term, and it leaves the perceived Gaussian variance undistorted. For the Bernoulli component, the same likelihood-ratio distortion of a probability p against a comparison $\bar{p}$ gives

$$p^\theta = \frac{p(p/\bar{p})^\theta}{p(p/\bar{p})^\theta + (1-p)((1-p)/(1-\bar{p}))^\theta} = p + \theta(p - \bar{p}) + O((p - \bar{p})^2);$$

routing the distortion through the hazard state, $\hat{p} = p(\hat{H})$ with $\hat{H}$ given by (D.1), reproduces this to first order in the news while keeping a single distortion parameter.

The impulse responses use a one-time unanticipated-shock convention. The baseline is the deterministic steady state of (D.1) with $D_t = 0$ and $\eta_t = 0$, around which the system is linearised; disaster exposure is priced through the certainty equivalent, while realisations are zero along the baseline. An unanticipated disaster sets $D_0 = 1$, and the response of any variable is the difference between the perturbed and the baseline path. Terms common to the two paths, including the calm-time news flow $-\theta \psi_H \bar{p}_t$ and all steady-state constants, cancel in the difference. Holding policies fixed, the perceived-

hazard-state impulse is $\partial \hat{H}_{t+1}/\partial D_0 = (1 + \theta)\psi_H \rho^t$; in the actual linearised system, the policy feedback $-\alpha \partial m_t/\partial D_0$ is incorporated into the transition matrix A_x of Proposition 3.

Equation (D.1) embeds diagnostic memory in the level of the hazard state estimate: the impact overshoot $\theta\psi_H$ is carried forward and decays at the persistence ρ of the perceived hazard state, rather than being re-benchmarked against the rational conditional each period. This is the parsimonious counterpart of distant memory, under which news fades from the comparison set only gradually and belief distortions persist, as in the business-cycle treatment of Bianchi, Ilut, and Saijo (2024) and the linear-model framework of L’Huillier, Singh, and Yoo (2024). The same convention governs the comparison probability: benchmarking news against the rolling one-step conditional $p(\hat{H}_t)$ instead of the distant-memory norm $\bar{p}_t$ would add a term $-\theta\psi_H p'(\hat{H}_t)$ to the perceived persistence, a third-order correction in $(\theta, \psi_H, \kappa_1)$ that leaves all results intact. Explicitly, the steady-state eigenvalue of the perceived hazard state becomes $\tilde{\rho} \equiv \rho - \theta\psi_H p'(H^*) < \rho$: the impact response is unchanged at $(1 + \theta)\psi_H$, the signal kernel becomes $(1 + \theta)\psi_H \tilde{\rho}^s$, and the cumulative signal weight falls from $(1 + \theta)\psi_H/(1 - \beta\rho)$ to $(1 + \theta)\psi_H/(1 - \beta\tilde{\rho})$; the capital/output block is unchanged except through policy feedbacks, so wherever the resource-defocusing channel dominates the diagnostic focusing channel in the medium run under the distant-memory convention, it does so a fortiori under the rolling one. Under the strict one-period-reversal convention, in which each period’s beliefs equal the rational conditional plus θ times current news only, the amplification applies only at impact and the signal kernel becomes $\psi_H \rho^s + \theta\psi_H \mathbf{1}\{s = 0\}$; all results go through, with the diagnostic focusing channel reverting faster, which strengthens the medium-run dominance of the resource-defocusing channel and hence the persistence of defocusing.

D.2 Steady-State Restrictions

Write the pre-realisation value as the certainty equivalent over the current disaster state,

$$V(K_t, H_t) = -\frac{1}{\Theta} \ln \left[(1 - p(H_t)) e^{-\Theta \tilde{V}(K_t, H_t, 0)} + p(H_t) e^{-\Theta \tilde{V}(K_t, H_t, 1)} \right],$$

where $\tilde{V}(K_t, H_t, D_t) = \max_{m_t, r_t} \{u(C_t) + \beta \mathcal{R}_t[V(K_{t+1}, H_{t+1})]\}$ is the post-realisation problem, the Bellman display in the text. Equations (26)–(27) are the steady-state first-order condition for r and the envelope condition for K along the no-disaster path; the tilted-mean corrections across the two period-ahead disaster branches are of order $p^* \delta$, and the Gaussian-tilt term from η is of order $\Theta \sigma_\eta^2$, both absorbed by the linearisation. Combining the two restrictions eliminates $u'(C^*)$ and yields $\beta(A + 1 - d) = 1$ up to the same corrections; with AK technology this is the existence condition for the deterministic interior steady state discussed in the text. The steady state additionally satisfies the mitigation first-order condition $(1 - \xi)V_K^* = \alpha|V_H^*|$ up to $O(p^* \delta + \Theta \sigma_\eta^2)$, which is required for $m^* > 0$. For (28), differentiate the outer certainty equivalent with respect to H_t : the hazard state enters through the current disaster probability and through both continuation branches, so

$$V_H = -p'(H) \mathcal{M}_V + (1 - \pi_V) \tilde{V}_H^n + \pi_V \tilde{V}_H^d,$$

where π_V and $\mathcal{M}_V$ are the tilted probability and the marginal value of probability reduction evaluated at the continuation-value gap $\Delta V = \tilde{V}^n - \tilde{V}^d$. By the envelope theorem applied to the post-realisation problem, and using $\partial \hat{H}_{t+1} / \partial \hat{H}_t = \rho$ under the distant-memory comparison, $\tilde{V}_H = \beta \rho V'_H$ in each branch, the cross-branch difference being of order δ . At the steady state this gives the fixed point $V_H^*(1 - \beta \rho) = -p'(H^*) \mathcal{M}_V^*$, which is (28), up to the $O(\delta)$ cross-branch variation just noted: the common-derivative substitution commits an error of order $\pi_V \delta$, second order in the small quantities and of the same status as the $O(p^* \delta)$ corrections to (26)–(27). Finally, for ΔV^* , evaluate the two post-realisation objectives at a common policy pair (m, r) , feasible in both branches for small δ by (23); their difference is

$$\Xi(m, r) = u(AK^* - m - r) - u(AK^*(1 - \delta) - m - r) + \beta \left\{ \mathcal{R}_0[V(K', \hat{H}')] - \mathcal{R}_1[V(K', \hat{H}')] \right\},$$

where $\mathcal{R}_D$ is the certainty equivalent under the branch- D conditional law, K' is common to the branches by (24), and the two laws differ only in their mean, by $(1 + \theta)\psi_H$, by (D.1). Optimality in each branch gives the sandwich $\Xi(m^{d*}, r^{d*}) \leq \tilde{V}^n - \tilde{V}^d \leq \Xi(m^{n*}, r^{n*})$; the branch optimisers differ by $O(\delta)$ and $\nabla \Xi = O(\delta)$, so the bounds agree to $O(\delta^2)$ and either may be expanded. The consumption term is $u'(C^*) \delta AK^* + O(\delta^2)$; the continuation term is β times the mean shift $(1 + \theta)\psi_H$ times the tilted-mean marginal continuation value $-V_H^* = |V_H^*|$, with curvature corrections of order $\psi_H^2 |V_{HH}^*| = O(\psi_H^2 \delta)$. Hence $\Delta V^* = \delta AK^* u'(C^*) + \beta(1 + \theta)\psi_H |V_H^*| + O(\delta^2 + \psi_H^2 \delta)$, and $\mathcal{M}_V^* \rightarrow \Delta V^*$ as $\Theta \rightarrow 0$, delivering the stated limits. $\square$

D.3 Two-Mode Persistence

For $t \geq 1$ write $IRF_m(t) = a\lambda_H^t - b\lambda_K^t$ with $a, b > 0$ and $\lambda_H, \lambda_K \in (0, 1)$; at $t = 0$ the representation identifies $g_D = a - b$, so mitigation falls on impact iff $b > a$. (i) If $b \geq a$ and $\lambda_K \geq \lambda_H$, then $b\lambda_K^t \geq a\lambda_H^t$ for every $t \geq 0$, so $IRF_m(t) \leq 0$ throughout. (ii) If $a > b$ and $\lambda_K > \lambda_H$, then $IRF_m(t) < 0 \iff (\lambda_K/\lambda_H)^t > a/b \iff t > t^* \equiv \ln(a/b)/\ln(\lambda_K/\lambda_H)$, and t^* is strictly decreasing in λ_K/λ_H : the defocusing spell $[t^*, \infty)$ starts earlier as λ_K/λ_H rises. (iii) At every fixed $t \geq 1$, $\partial IRF_m(t)/\partial \lambda_K = -bt\lambda_K^{t-1} < 0$: a larger capital eigenvalue deepens the shortfall pointwise. (iv) $\sum_{t \geq 0} IRF_m(t) = a/(1 - \lambda_H) - b/(1 - \lambda_K)$ is strictly decreasing in λ_K and increasing in λ_H . Statements (ii)–(iv) are the formal content of the claim that the depth and duration of defocusing rise with λ_K/λ_H . $\square$

D.4 Proof of Proposition 3

No embedding is required here: within the linearisation the policy rules are linear in the shock vector, so the response to the binary disaster coincides with the difference between the $D_0 = 1$ and $D_0 = 0$ paths, and $IRF_m(t)$ denotes this difference throughout. Substitute the local policy rules into the capital law (24) and the perceived-hazard-state law (D.1), after subtracting the steady-state constants and setting

$\eta_t = 0$ for the deterministic MIT impulse response. This gives

$$\tilde{K}_{t+1} = (1 - d + h_K + \xi g_K)\tilde{K}_t + (h_H + \xi g_H)\tilde{H}_t + (h_D + \xi g_D)D_t,$$

and

$$\tilde{H}_{t+1} = -\alpha g_K \tilde{K}_t + (\rho - \alpha g_H)\tilde{H}_t + ((1 + \theta)\psi_H - \alpha g_D)D_t,$$

which is the stated system $x_{t+1} = A_x x_t + b_D D_t$ with A_x and b_D as displayed in the proposition. With $x_0 = 0$, $D_0 = 1$, and $D_t = 0$ for $t \geq 1$, iterating gives $x_1 = b_D$ and, for $t \geq 1$, $x_t = A_x^{t-1} b_D$. Substituting into the mitigation rule $\tilde{m}_t = g_K \tilde{K}_t + g_H \tilde{H}_t + g_D D_t$ yields $IRF_m(0) = g_D$ and $IRF_m(t) = \begin{pmatrix} g_K & g_H \end{pmatrix} A_x^{t-1} b_D$ for $t \geq 1$. Summing over $t = 0, \dots, T$,

$$\sum_{t=0}^T IRF_m(t) = g_D + \begin{pmatrix} g_K & g_H \end{pmatrix} \left(\sum_{j=0}^{T-1} A_x^j \right) b_D,$$

and when $I - A_x$ is invertible the geometric sum is $\sum_{j=0}^{T-1} A_x^j = (I - A_x^T)(I - A_x)^{-1}$, delivering the stated cumulative expression. $\square$