

Adapting to Physical Climate Risk: Evidence, Decisions and the Missing Bridge to Finance

*Iñigo J. Losada **

Abstract

Physical climate risk has crossed from projection to observation. 2024 was the first calendar year more than 1.5°C above pre-industrial levels, 2015–2025 were the eleven warmest years on record, and the rate of global sea-level rise has doubled since satellite measurements began. Yet the argument of this paper, developed from the author's participation in the panel “Transition Plans and Adaptation to Physical Risk” at the Banco de España High-Level Conference on Climate Risk and Sustainable Finance (Madrid, 25 March 2026), is that physical climate risk is now as much a matter of decisions as of evidence. The hazard is measured and rising; how much of it becomes risk depends on adaptation, that is, on choices about what to protect, how much, for whom and when. Those choices are taken at local scale, under deep uncertainty and against multiple criteria, and the financial system has yet to internalize any of the three.

Four messages follow. First, what matters for risk is the shift in extremes, their sequencing and compounding, and the erosion of resilience, the function a system keeps and the speed at which it recovers, when events arrive faster than recovery allows; the transmission from hazard to sectoral damage to macro-financial impact remains the weakest link in the chain. Second, risk is in part a construct of choices (what to protect, how much, for whom and over what horizon), which is why adaptation has no equivalent of the 1.5°C target and must instead be governed through tolerable residual risk and flexible adaptation pathways set at local scale. Third, the implementation gap is a gap of governance and of funding permeability across central administration, regional, municipal and private levels, not a knowledge gap alone. Fourth, “optimal” adaptation cannot be reduced to minimizing the public cost or achieving insurability, these are constraints on a multi-criteria decision, not its objective, and the adaptation effort itself, if compressed in time or concentrated in indebted territories, is a source of financial risk in its own right. The paper closes with what physical-risk science can and cannot yet deliver to prudential transition plans, stress tests and the forthcoming European climate-resilience framework. Adaptation must become legible to finance, through physically consistent scenarios calibrated on well-documented local events; and finance would do well to stop treating it as an exogenous parameter.

Keywords: physical climate risk, adaptation, adaptation pathways, coastal risk, financial stability, transition plans.

1. INTRODUCTION: CONTEXT AND SCOPE

On 25 March 2026 the Banco de España hosted, with the Network for Greening the Financial System (NGFS), the Centre for Economic Transition Expertise (CETEx) and the TPI Global Climate Transition Centre at the London School of Economics, the High-Level Conference “Climate Risk and Sustainable Finance: Bridging Policy, Research and

* Distinguished Researcher. Instituto de Hidráulica Ambiental-IHCantabria, Universidad de Cantabria.

Practice”. Its three sessions (climate data and analytical tools, transition plans and adaptation to physical risk, and the greening of central-bank portfolios) were framed by an explicit choice: while mitigation remains essential, the day would place particular emphasis on adaptation, as the dimension now critical for protecting communities and for the stability of economic and financial systems. The conference thus marked a visible shift in the central-banking climate agenda, from transition risk, where the analytical machinery has a decade of development behind it, to physical risk and adaptation, where it is still being built.

This paper takes up one strand of that agenda. It is written from the author’s participation in the panel on transition plans and adaptation to physical risk and is organized around the three questions the panel was asked to address. What does the evidence say about the impacts of physical risk, and where is adaptation most useful and most implementable now? Is an overarching adaptation target conceivable, as the temperature goal is for mitigation, and where are the largest gaps between what could be done and what is being done? And, from the floor, can the optimal level of adaptation be defined as the level that minimizes the cost to government and makes the residual risk insurable? A fourth theme runs beneath all three and is taken up in Section 6: the difference between the partial-equilibrium view of an institution assessing its own exposures and the general-equilibrium view that a public authority must hold, between private cost and social cost.

This paper takes the adaptation strand and tries to turn the panel’s insights into a structured argument. Its spine is a single claim: physical climate risk is now as much a matter of decisions as of evidence. The hazard is measured and rising, and the evidence on that is settled; how much of it becomes risk depends on adaptation, that is, on choices about what to protect, how much, for whom and when. Those choices have three features that distinguish them from mitigation and that the financial system has yet to internalize: they are taken at local scale, under deep uncertainty and against multiple criteria. Adaptation is likely to remain the least developed part of climate policy, and the least legible to finance, until its governance and its metrics reflect those three features. Section 2 sets out the evidence and why the signal lives in the extremes. Section 3 presents the adaptation cycle as a sequence of decisions rather than a modelling exercise. Section 4 explains why adaptation has no 1.5°C and what governs it instead. Section 5 addresses the question of optimal adaptation and insurability. Section 6 draws the bridge to the supervisory and financial agenda, and Section 7 concludes.

2. THE EVIDENCE: FROM MEANS TO EXTREMES, FROM EVENTS TO SEQUENCES

The scientific evidence on physical climate risk is now overwhelming, and the signal that matters for risk does not live in the means. It lives in the extremes, in their sequencing, and

in their compounding across hazards and regions. Decades of improved observation converge on one conclusion: hazards are increasing in severity and in frequency, in their averages and in their tails, for almost every climate variable one chooses to examine.

2.1. The observed record

The World Meteorological Organization's State of the Global Climate 2024 confirmed that 2024 was the warmest year in the 175-year observational record, at $1.55 \pm 0.13^\circ\text{C}$ above the 1850–1900 average, and likely the first calendar year above 1.5°C ([WMO, 2025](#)). One year above the threshold does not mean the Paris goal has been breached (the long-term warming level was estimated at $1.34\text{--}1.41^\circ\text{C}$) but the direction is unambiguous. The 2025 edition, released on 23 March 2026, two days before the conference, added that 2015–2025 were the eleven warmest years on record, that 2025 was the second or third warmest at about 1.43°C , that the ocean continues to absorb around 90% of the excess heat, and that the Earth's energy imbalance, tracked as a headline indicator for the first time, reached a record high ([WMO, 2026](#)).

For a coastal scientist, like this author, the single most consequential number in these reports is the rate of sea-level rise. Global mean sea level rose at 2.1 mm per year over 1993–2002 and at 4.7 mm per year over 2015–2024: the rate has doubled within the satellite era ([WMO, 2025](#)). Sea level is now about 11 cm above its 1993 level, and the rate since 2012 exceeds that of 1993–2011 ([WMO, 2026](#)). Sea level is the baseline against which every coastal return period is computed. Shift the baseline and the whole distribution of extreme water levels shifts with it: the IPCC's Sixth Assessment projects that extreme sea-level events that occurred once per century in the recent past will occur at least annually at more than half of the world's tide-gauge locations by 2100 (IPCC, 2021). In plain terms, under some scenarios the event that engineers designed against as the hundred-year storm becomes an annual occurrence. A design standard becomes a season.

The year in which this paper is written has extended the record rather than interrupted it. August 2026 was the warmest August in the ERA5 dataset (Hersbach et al., 2020) and, tied with July 2023, the joint warmest month ever recorded, at 1.65°C above pre-industrial levels; it was the first month above 1.5°C since November 2025, and June–August 2026 was the joint warmest boreal summer on record, level with 2024. Sea-surface temperature over the extra-polar ocean reached an August record of 21.07°C , the joint highest for any calendar month, with strong El Niño conditions developing in the equatorial Pacific ([C3S, 2026a](#)). For western Europe the summer of 2026 was the warmest on record, surpassing 2003, at 21.69°C , 2.54°C above the 1991–2020 average, marked by a succession of exceptionally early, persistent and intense heat waves from the second half of May, widespread drought, and exceptionally low flows on the Rhine and the Danube ([C3S, 2026b](#); [C3S, 2026c](#)).

Spain sits at the center of that anomaly. AEMET's summer balance, published on 22 September 2026, records the warmest summer since the official series began in 1961 and, on a reconstructed series, since at least 1869, with a peninsular mean of 24.5°C, 2.4°C above the 1991–2020 reference and 0.3°C above the previous record, set only a year earlier. Sixty-one days were spent under heat-wave conditions, against the previous record of 41 in 2022; July, at 25.7°C, tied July 2022 as the warmest month of the series; summer rainfall was 39.6 mm, 56% of normal and the third lowest since 1961; and fully summer conditions extended over four months, from an episode 5°C above normal in the second half of May to a fifth heat wave in the first week of September ([AEMET, 2026a](#); [AEMET, 2026b](#)). The ten warmest summers of the Spanish series all belong to this century, and seven of the eight warmest to the years since 2015. Heat and drought translated into fire: by early September close to 296,000 hectares had burned in Spain according to the European Forest Fire Information System, following the 393,000 hectares of 2025, the worst year in the recent record (EFFIS, 2026). None of these numbers is a projection. They are the observed record of the year in which this paper was written, and they sit where the trend would have placed them.

2.2. Frequency, sequencing and the erosion of resilience

It is tempting to read the evidence as a story about intensity. The story that matters more for risk is about frequency and sequencing, and the concept that captures it is resilience. In the IPCC's definition, resilience is the capacity of a system (social, economic or ecological, and usually all three at once) to cope with a hazardous event, trend or disturbance, responding or reorganizing in ways that maintain its essential function, identity and structure while retaining the capacity for adaptation (IPCC, 2022). Engineers make the idea operational as a curve: the function lost when the event strikes and the time taken to recover it (Bruneau et al., 2003). Seen that way, resilience has three components that the adaptation cycle of Section 3 must treat separately: the capacity to absorb a shock, the capacity to recover from it, and the capacity to adapt so that the next shock does less harm. Adaptive capacity is the third of these, the forward-looking one; it is a component of resilience, not a synonym for it.

Every extreme event consumes part of all three. Physical defenses are damaged, budgets are drawn down, institutions are exhausted, households and firms lose reserves, and the attention that adaptation requires goes to reconstruction instead. Recovery takes time. If the next event arrives before recovery is complete, the risk it poses is higher than the hazard alone would suggest, because the system is still weakened. Its vulnerability has risen and its resilience has fallen in the interval. Risk increases not because the storm is stronger but because the system has not yet recovered from the last one.

The 2017 Atlantic hurricane season was the first in which the United States and its territories were struck by three Category 4 landfalls (Harvey, Irma and Maria) in a single season. Puerto

Rico, weakened by Irma, was devastated by Maria two weeks later; the second event did far more harm than it would have done to an intact system. We now observe the same pattern across heat waves, droughts, wildfires, extreme precipitation, river flooding and coastal flooding. Southern Europe's summers since 2022 have been a sequence of this kind.

The Iberian winter of 2026 supplied a textbook case of sequencing weeks before the conference. Between 22 January and 14 February, seven named Atlantic storms (Ingrid, Joseph, Kristin, Leonardo, Marta, Nils and Oriana) crossed the peninsula in succession with no recovery periods between episodes; soils saturated progressively and the hydrological response intensified with each storm. Leonardo, fed by a subtropical atmospheric river, left 577 mm in 24 hours at Grazalema (Cádiz), a national record, and more than 2,500 mm fell there in little over three weeks; AEMET ranks the sequence among the most significant since at least 1961 and among the three largest of this century for accumulations over periods of two weeks to a month ([AEMET, 2026c](#)). The same storm train struck Portugal and Morocco, and a rapid attribution study concluded that climate change had made the rainfall more intense. The Consorcio de Compensación de Seguros treated the period 20 January to 20 February as a single episode of atypical cyclonic storm comprising eight storms, and activated its claims-management agreement with the insurers' association on account of the episode's intensity, duration and territorial extent ([CCS, 2026a](#)). The damage done by the eighth storm was damage to a system whose resilience the seventh had already spent.

A second feature is compounding: events that combine across hazards (drought followed by flood), across regions (the same storm system hitting several countries), or across time (back-to-back seasons). Italy's recent record makes the statistical point (figures from Battiston, 2026). Direct costs of climate-related disasters average about €3 billion per year, but the average is made up of many years near €1 billion and a few years at €10 billion or more; a handful of events account for 60–80% of cumulative damage. The 2022 drought in the Po valley, at roughly €6 billion, was followed within a year by the 2023 floods in Emilia-Romagna at roughly €10 billion. The distribution of losses is not Gaussian; it is closer to a generalized Pareto, and its tail is where the decision-relevant information sits. Please note that the generalized Pareto distribution (GPD) is a family of continuous probability distributions used mainly to model the data points that exceed a high threshold. Spain's 2026, a record storm train in winter, a record heat and fire season in summer, is compounding across time within a single year.

2.3. Southern Europe and Spain as a hotspot

The projections are as clear as the observations, and for Spain, though one could say the same of Portugal, Greece or Italy, they are unfavorable. The IPCC's Sixth Assessment Report identifies the western Mediterranean as a hotspot in which water scarcity, more frequent hydroclimatic drought, intensifying wildfire and the growing vulnerability of densely populated

coasts coincide (IPCC, 2022). Spain's first integrated national assessment, the *Evaluación de Riesgos e Impactos derivados del Cambio Climático en España* (ERICC-2025), mandated by Law 7/2021, characterized those risks sector by sector, including, for the first time, the financial and insurance sector and the cascading effects that originate in the direct impacts on other sectors, and concluded that, for all the regulatory and analytical progress made, the tools available remain insufficient to represent the complexity of physical risks and their interaction with macroeconomic and financial risk (Losada et al., 2025). The first European Climate Risk Assessment (EEA, 2024), which identified 36 major climate risks and singled out southern Europe as a hotspot, reached the same conclusion at continental scale. Physical risk is material now and will become more so; how much of it becomes financial risk is the question the rest of this paper addresses.

2.4. The weakest link

What the evidence does not yet give us is the chain that connects a physical event to a balance sheet. Between a hazard and a financial exposure lie several links: physical damage to assets, disruption of sectoral activity, macroeconomic shock, transmission to credit, market and operational risk. Each with its own data, models and uncertainties, and each studied by a different community. The climate stress tests run so far have, understandably, bridged that chain by assumption. The shocks applied to the macroeconomic models (i.e. to total factor productivity, sectoral output, relative prices or sectoral income) are specified consistently with NGFS or ECB scenarios, but, to the author's knowledge, they are not derived from any physical reconstruction of a plausible event. Such exercises give a valuable first view of financial vulnerability, but their limits should be stated plainly. They cannot distinguish a severe coastal flood from a prolonged drought, they do not reproduce the intensity or duration of the phenomenon, they offer no test of coherence between the shock applied and the physical damage that could plausibly produce it, and they leave the transmission mechanisms uncalibrated against evidence. Building the link *physical event* → *sectoral damage* → *macroeconomic shock* → *financial impact* on physical rather than assumed foundations is the frontier on which physical scientists and financial economists have to meet. The rest of this paper is about what the physical-risk side can bring to that meeting, and what it cannot yet.

3. RISK IS A DECISION: THE ADAPTATION CYCLE

Adaptation is arguably the least developed part of climate policy, in part because it forces a sequence of decisions that mitigation does not. Mitigation has one metric, one direction and, in principle, one global objective. Adaptation asks, before any measure is built: what system, what risk, what goal, what residual risk, what strategy, and when. The adaptation cycle is usually drawn as a loop of assessment, planning, implementation and monitoring. In this

context the author prefers to present it as five decisions, because that is what it amounts to in practice, and because each decision is a place where the physical-risk community and the financial community still tend to talk past each other.

Figure 1

The adaptation cycle as five decisions



Monitoring feeds back into the assessment, which is what makes pathways possible.

Fuente: Elaboración propia.

3.1. Decision one: define the system

The first decision is what we are adapting. Risk, in the IPCC's formulation, is a function of hazard, exposure and vulnerability. Hazard is what the climate does; exposure is what stands in its way; vulnerability is how badly what stands there is hurt. Before any of these can be quantified one must draw a boundary: which elements are exposed, what their vulnerability is, and, the part most often neglected, which processes connect them. A port, the road that serves it, the industrial estate behind it and the housing on the reclaimed marsh are one system, not four assets. The boundary determines every number downstream, and it is a choice. Resilience is a property of that system, not of any of its assets; the port with a second road recovers in days; the port without it recovers when the road does.

3.2. Decision two: assess risk honestly

The second decision is how to assess risk, and here the physical-risk community owes the financial community some candor about its own products. Take flood modelling, since flooding (i.e. river, pluvial and coastal) is the hazard most often mapped and most often bought. A flood map depends on the model that produced it. A process-based hydrodynamic model and a simplified planar-surface model can give very different inundation extents for the same event. It depends on the topography beneath it: a coarse or inaccurate digital elevation model will overestimate flooding. And it depends on whether existing protection and existing resilience (i.e. the dike that is already there, the drainage that was upgraded in 2015, the early-warning system that empties the ground floor before the surge arrives, the recovery capacity of the institutions and the economy behind them) have been taken into account. Leave them out and the map will overstate risk, often by a large margin.

The consequence is that risk assessments are not comparable unless the approach is coherent across the whole chain. Two hazard providers can return different exposures for the same mortgage book, and both can be internally consistent. Establishing the protection and the resilience that already exist, up front, is a precondition for knowing where risk actually stands today, and today's risk is the reference point for every adaptation goal.

3.3. Decision three: set the goal

The third decision is the one that has no analogue in mitigation. We know what we want from mitigation: net-zero emissions on a given date, consistent with a temperature limit. What do we want from adaptation? One must decide what to adapt, how much, for what time horizon and for what climate scenario. Each of these is a choice, and the choices differ from one location to the next, because adaptation is implemented at local scale. The variable that ultimately must be fixed is the level of residual risk a community is willing to tolerate; the risk that remains after the measures are in place. Every adaptation plan the author has been involved in has reached the point where this question had to be put to the people who would have to answer for it, and every time it has been the hardest decision on the table. It is also the decision that a bank, an insurer or a supervisor most needs to know, because tolerable residual risk is, in financial language, the risk appetite of the territory in which the asset sits. Increasingly the goal is also stated in terms of resilience: how much function a system must keep during the event and how quickly it must recover (hours for a hospital, days for a port, a season for a beach). The two formulations are complementary. Residual risk bounds the loss, resilience bounds the disruption, and a lender or an insurer needs both.

3.4. Decision four: choose the strategy

The fourth decision is the strategy. Whatever the hazard, an adaptation measure can act on one of five things: 1) the hazard that reaches the asset, 2) the exposure that stands in its way, 3) the vulnerability of what is exposed, 4) the resilience with which the system keeps and recovers its function, or 5) the allocation of the residual risk. Coastal planners know the first four as protect, retreat, accommodate and prepare, and the fifth as insurance; the same logic governs river floods, drought, heat and fire, and it maps onto the vocabulary supervisors are beginning to use (i.e. invest, avoid, accept, transfer) a correspondence the author returns to in Section 6. Each strategy can be delivered by very different combinations of measures, grey and green, structural and operational, and the combination is itself a design choice.

Table 1

The strategy space across hazards and its risk-management analogue

Strategy	What it changes	Examples across hazards	Risk-management analogue
Reduce the hazard	The event that reaches the asset.	Dikes, dunes and wetlands against flooding; storage and reuse against drought; urban shade and green cover against heat; fuel management against fire.	Invest: lowers the probability of loss.
Reduce exposure	What stands in the way.	Setback lines and managed retreat; land-use planning in floodplains; crop and land-use change in water-stressed basins; relocation of activity.	Avoid: lowers the exposure at loss.
Reduce vulnerability	The damage per unit of hazard.	Flood-proofing and elevation; early warning; irrigation modernization and efficiency; building design and heat-health plans.	Invest: lowers the loss given event; the asset stays, its damageability falls.
Strengthen resilience	The function kept during the event and the speed of recovery.	Redundancy in transport, power and water networks; emergency response and business-continuity planning; contingency finance and social protection; ecosystem restoration.	Invest: shortens the duration of loss (business interruption, recovery time).
Transfer the residual	Who pays.	Insurance and public compensation schemes; agricultural insurance; parametric covers.	Transfer: moves what remains to a carrier able to hold it.

Note: “accept” is not a lever. It is the tolerable residual risk of Section 3.3 that remains once the other five have been applied.

3.5. Decision five: sequence under deep uncertainty

The fifth decision is when to act, and it is governed by uncertainty. The modelling chain runs from emissions scenarios, whose probabilities we do not know, through the downscaling of

global climate projections to regional and local hazard, to exposure, to vulnerability, to risk. At every link uncertainty enters and must be tracked, not averaged away. For some links the uncertainty is deep. The IPCC's Sixth Assessment could not rule out global mean sea-level rise approaching 2 m by 2100 under a very high emissions scenario, because of ice-sheet processes whose physics remain poorly constrained (IPCC, 2021). No single design horizon can accommodate that spread.

The sensible answer is not to pick the solution we believe will hold forever. It is to combine implementation with observation: to define a set of actions, each with the conditions under which it stops being adequate (i.e. its adaptation tipping point) and to monitor how risk evolves so that the next intervention is triggered when the observed system, not the projected one, requires it. Recovery times after each event are the most direct evidence of where a system stands on its pathway; resilience is not only a design objective but the variable being monitored. This is the logic of dynamic adaptive policy pathways (Haasnoot et al., 2013; Fernandez-Perez et al. 2025), which has become the standard framing for long-horizon adaptation planning, from estuaries and deltas to urban water supply. Pathways buy flexibility and avoid lock-in, at the cost of an institutional obligation that few administrations have so far taken on: someone has to keep watching. The moment one tries to connect this decision structure to economic and financial impacts, things become harder still, because the financial system prices assets against a fixed horizon while the physical system is being managed against a moving one.

4. WHY ADAPTATION HAS NO 1.5°C: GOALS, SCALE AND GOVERNANCE

There is no global adaptation target comparable to 1.5°C, and it is hard to see how there could be one. What exists instead is an indicator framework, a finance pledge, and a widening gap between the assessments we have produced and the adaptation we have implemented. That is the answer to the panel's second question, and the months since the conference have, if anything, sharpened it.

4.1. The global stocktake and Belém

The global stocktake concluded at COP28 (2023) in Dubai and the process that followed under the Global Goal on Adaptation taught us something uncomfortable: the world has produced a great many climate risk assessments (one can argue about their quality, but not about their number) and, measured against stated objectives, has implemented very little adaptation (see for example, Magnan et al., 2023). COP30 in Belém (November 2025) was billed as the adaptation COP. Its outcome was a list of 59 voluntary indicators, the Belém Adaptation Indicators, covering the eleven targets of the UAE Framework for Global Climate

Resilience and reduced from an initial pool of roughly ten thousand candidates. On finance, the decision calls on parties to at least triple adaptation finance by 2035, without a reference year. Against that pledge stands the arithmetic of the UNEP Adaptation Gap Report 2025: developing countries will need \$310–365 billion a year by 2035, while developed countries provided \$26 billion in 2023, down from \$28 billion the year before, a gap of twelve to fourteen times current flows, with the Glasgow goal of doubling adaptation finance by 2025 on course to be missed ([UNEP, 2025](#)).

An indicator framework is not a target. Fifty-nine indicators describe progress; they do not say how much progress is enough, because that question can only be answered system by system, community by community. This is less a failure of negotiation than the nature of the problem.

4.2. Local risk, national plans

The deepest reason there is no 1.5°C for adaptation is scale. Adaptation is implemented locally, but adaptation planning is national. A national plan can set direction, allocate resources and define methods; it cannot decide the residual risk that a particular estuary town, irrigation district or heat-exposed city will tolerate. Spain's coastal adaptation strategy, mandated by Law 2/2013 and approved on 24 July 2017 on the basis of the national risk assessments produced from 2014 (MITECO, 2017 and Losada et al., 2019), illustrates the tension. Its overarching objective, to keep coastal risk at the level assessed at the time, under rising hazard, is a legitimate and communicable national goal; operationalizing it means going stretch by stretch, equipping each with a defensible risk metric and a portfolio of measures that keeps risk below the threshold as sea level and storminess rise. Drought plans, flood-risk management plans and heat-health plans face the same arithmetic: the national goal is one sentence, the local plans are thousands of decisions, most of them still to be taken.

4.3. The two gaps

Two gaps stand out between what could be done and what is being done.

The first is methodological and institutional at once. We do not yet know how to build the flexible adaptation pathways of Section 3.5 as living instruments. We know how to draw them; we do not know how to fund, staff and govern the monitoring that gives them meaning. A pathway without a monitoring programme is a plan with an expiry date that no one is watching.

The second is governance and the permeability of funding. National funds may exist, but the channels through which they reach a municipal drainage scheme or a regional dune-restoration

programme are narrow and slow. In Spain the chain runs from the central administration to the autonomous communities to the municipalities, and from there to the private owners of most of the exposed assets, and each level may set its adaptation goals differently and hold interests that do not align. Every measure implemented at one level interacts with the system in which it is embedded: a seawall built by one municipality changes erosion in the next; a retreat decided by a regional plan changes the value of assets financed by a national bank. Adaptation governance is therefore not a matter of assigning responsibility but of managing interactions, and this, at least as much as the science, is where implementation tends to stall.

4.4. The European turn

The European Union is now attempting what national plans have so far struggled to achieve: to move adaptation from strategy to obligation. The first European Climate Risk Assessment (March 2024) concluded that the Union is not sufficiently prepared and that some risks are already at critical levels, and the European Court of Auditors reached similar conclusions in 2024. The Commission is preparing the European Climate Resilience and Risk Management Integrated Framework, the package announced in the 2024 political guidelines as the European Climate Adaptation Plan, with legislative and non-legislative elements, expected in the second half of 2026 ([European Commission](#)). The consultation that preceded it called for resilience-by-design in investment, harmonised climate risk assessments and stable long-term adaptation financing. If the framework delivers harmonised assessment methods and binding resilience requirements for climate-vulnerable investment, it will for the first time give supervisors a policy anchor for adaptation comparable to those they have for mitigation. Whether it does is, at the time of writing, open.

Seen from the balance sheet, the same problem appears in a different vocabulary. The supervisory diagnosis presented in the panel is that prudential transition plans remain mitigation-centric, that the systematic integration of adaptation is limited, and that the lack of relevant metrics and targets is the binding constraint, because adaptation is local and context-specific, without a single global quantified objective, and because quantitative policy anchors are few (Casillas, 2026).

5. BEYOND COST: MULTI-CRITERIA DECISIONS, RESIDUAL RISK AND INSURABILITY

The optimal level of adaptation should not be defined as the level that minimizes the cost to government and makes the residual risk insurable. Cost and insurability are constraints on the decision; they are not its objective function. The question was put from the floor, and it deserves more space than the panel could give it, because the proposition is understandably attractive to the audience this paper is written for.

5.1. Why the proposition is tempting, and where it falls short

The proposition is tempting because it would close the problem. If adaptation is the quantity that minimizes public expenditure subject to the constraint that what remains can be priced and transferred to an insurance market, then the goal of Section 3.3 has a formula, the governance problem of Section 4.3 reduces to budgeting, and the financial system knows where it stands. However, it runs into three difficulties.

First, adaptation is not only a cost. It has social and environmental implications, co-benefits and collateral effects, and all of them have to be accounted for. A dune-restoration programme protects a beach and also restores a habitat, sustains a tourism economy and changes who can live where; a seawall protects assets and also may contribute to destroying a beach, alters sediment transport downstream and can lock a community into a defended position it may later regret. This is why adaptation plans generally select their measures through multi-criteria analysis rather than through cost-benefit analysis alone.

Second, the criterion that decides is frequently not the one the economist expects. After the 2011 tsunami, some communities on the Tōhoku coast rejected the seawall that engineering and cost analysis recommended, while others rejected relocation inland and demanded a seawall in order to stay where their families had lived for centuries. Neither was irrational; both were optimizing an objective function that did not appear in the cost model, and something similar can be expected in many of the floodplains and irrigated districts in Spain from which a plan may propose to withdraw.

Third, insurability is a property of the residual risk, and the residual risk is fixed by the tolerable-risk decision of Section 3.3, which belongs to the community, not to the insurer. Insurability matters, the protection gap is real, and a residual risk that cannot be transferred lands on public budgets, but a plan whose only purpose is to make the leftover risk insurable has inverted the order of decision. One decides what to protect and how much and then asks how the remainder can be shared.

5.2. Scale and who pays

There is a fourth reason, and it concerns scale. Adaptation will require very large investment. The author is entirely in favor of a climate-resilient development policy that combines mitigation and adaptation. But it is reasonably clear in advance that public budgets alone will not protect what is exposed: some 8,000 km of coastline, the floodplains on which much of Spain's industry and housing stands, the irrigated agriculture of the water-stressed south and east, and the cities in which heat now costs lives and productivity. Under current governance the public administration bears most of the responsibility for protection, and the plans are public documents. They will not be implementable without the private sector, i.e. the owners, developers, insurers and lenders whose assets are what is being protected, and who at present

treat protection as something the State provides. The figures presented for the EU in the conference keynote, annual adaptation needs of about €83 billion against current flows of €13 billion, a gap of €70 billion of which roughly €67 billion falls on the public sector, with private finance covering only around 5% of adaptation in high-income countries (Markandya, 2026), describe an arrangement that is unlikely to scale, and point precisely to the fiscal channel through which physical risk reaches sovereign creditworthiness.

5.3. The cost of resilience as a financial risk

That fiscal channel points to a dimension to which the adaptation debate, preoccupied with the cost of inaction, has paid comparatively little attention; the adaptation effort itself can generate macroeconomic and financial risk. Adaptation at the scale the evidence requires means a large volume of investment, concentrated in a few sectors (i.e. health, construction, water, energy, transport and territories) and in some cases needed within short periods. Investment of that kind has the familiar expansionary effects of infrastructure spending; it can also crowd out other uses of capital, generate cost inflation in civil engineering when it is compressed in time, raise public debt precisely at the regional and municipal levels where exposure is greatest and fiscal room smallest, and increase leverage in capital-intensive sectors that must adapt before their investments pay back. The risk is less one of insufficient adaptation than of disordered adaptation. An urgent need to reinforce defenses, modernize urban water systems or adapt critical infrastructure, arriving after the event and financed under pressure, can produce liquidity stress in indebted administrations and transitory deterioration in sectors that need credit at the moment their collateral has lost value.

A financial-stability framework should therefore be able to compare trajectories: one in which adaptation reduces physical risk but generates short-term financial tension; one in which under-investment accumulates physical risk and future losses; and one in which adaptation distributed sensibly over space and time reduces structural risk and stabilizes exposed sectors. The indicators are not exotic. Adaptation investment as a share of GDP, additional indebtedness attributable to adaptation and its sectoral distribution, the absorption thresholds beyond which tension appears, the net financial risk once the costs and the benefits of adaptation are both counted but, to the author's knowledge, no supervisory framework currently produces them. Anticipating the risks of action as well as of inaction is what it would mean to treat adaptation as a financial-stability question.

5.4. Adaptation and mitigation are not substitutes

A related question is how adaptation and mitigation interact, and the answer completes the argument. Mitigation is a global public good, which explains why it is undersupplied. Adaptation is a local public good, which is why it is undersupplied too, only at a smaller

scale. A firm can build a wall around its plant, but unless the municipality has built the dike upstream the wall protects nothing. And there is only so much one can adapt to. In the absence of mitigation, at some point no level of adaptation works; and the less we mitigate, and the later we adapt, the more expensive adaptation becomes. The two are complements in a climate-resilient development path, and a financial framework that treats adaptation as the residual after mitigation has been costed is likely to misprice both.

6. THE MISSING BRIDGE: MAKING ADAPTATION LEGIBLE TO FINANCE

The bridge between physical-risk science and finance is not a better hazard map. It is a shared language about decisions, tails and residual risk and a recognition on the financial side that adaptation is an endogenous variable, set by public decisions at local scale, rather than a parameter to be read off a data vendor's file. This section aims at drawing that bridge from both ends. From the supervisory agenda as it now stands, and from the physical-risk science that has to feed it.

6.1. What the supervisory agenda now requires

The regulatory trajectory is by now well defined. The ECB moved from setting expectations in its November 2020 guide on climate-related and environmental risks, through thematic reviews, to the first supervisory climate stress test in 2022; the banking package (CRR3 and CRD6) introduces a formal definition of ESG risk and requires prudential transition plans (risk-management instruments distinct from the disclosure-oriented plans of the CSRD) and the European Banking Authority's guidelines on the management of ESG risks bring these requirements into the Internal Capital Adequacy Assessment Process (ICAAP), risk appetite and strategy, including the quantification of short-, medium- and long-term financial impacts. The supervisory assessment of progress, as presented in the panel, is that it is real but uneven (Casillas, 2026): the regulatory perimeter has expanded faster than implementation; institutions still struggle with data, with integrating pathways into credit models and stress tests, and with governance; and adaptation in particular has barely entered the frame. The prescription for institutions is pragmatic. Scope and gather data, assess exposure and vulnerability, translate the assessment into a choice among avoid, accept, transfer and invest, and track progress along a maturity pathway without waiting for perfect data. The reader will recognize the first two steps as decisions one and two of Section 3, and the four options as the strategy space of Section 3.4.

6.2. Two conceptual blind spots

Two blind spots remain that no amount of data will remove, because they are conceptual (Battiston, 2026). The first is the reliance on expected values. Public decision-makers, and many financial ones, reason in terms of expected annual loss: a few tenths of a percent of GDP,

small next to a geopolitical shock, and therefore not a priority. But the distribution of climate losses is fat-tailed, a few events carry most of the damage, and the tail events are not one-in-a-million; they are events that will occur within the lifetime of the assets being financed. A decision framework that ignores the tail will, in this domain, produce the wrong decision. The NGFS short-term scenarios are the first set intended for financial use in which extreme events enter as return periods rather than averages, the equivalent of a value-at-risk on the loss distribution, although a 1-in-50 combined hazard is too extreme for some users and not extreme enough for insurers, whose reference is 1-in-100.

The second blind spot is indirect impact. What can be counted after a flood is the direct loss, often proxied by the aid claimed from the central administration. What is not counted, because it is confounded by everything else, is the long-run effect on the attractiveness of the territory for investment, on the confidence of people in staying, on the fabric of the local economy. Taking it into account changes the adaptation decision; ignoring it produces a different one. This is the distinction between partial and general equilibrium, between private and social cost, made concrete. An individual institution will, correctly, assess its own exposure. Someone must hold the general-equilibrium view, the indirect losses, the correlated exposures across institutions, the fiscal backstop and that someone is the public authority. Physical-risk science can supply the hazard and the direct-damage functions; it cannot, at present, supply the indirect-loss function, and it should say so.

On adaptation itself, one point deserves to be repeated to every scenario user. In a model, adaptation is an adjustment to the vulnerability function of the agents that are hit, and data on how much firms are actually adapting on the ground are very scarce. Results without adaptation are an upper bound. Results with adaptation are, at present, an assumption.

6.3. Adaptation as an endogenous variable

The physical risk on a bank's mortgage book in a floodplain, on an insurer's commercial-property portfolio in a water-stressed basin or on critical infrastructure depends on decisions the institution does not take: whether the municipality builds the dike, whether the basin plan restricts abstraction, whether the regional plan permits retreat, whether the national strategy funds the works. For the institution these decisions are externalities. They are, however, observable externalities, and the adaptation pathway of Section 3.5 is exactly the object a supervisor could ask to see (i.e. the sequence of planned actions, the residual risk each is designed to hold, the trigger conditions under which the next action is due, and the monitoring that watches for them). A territory with a published pathway, a stated tolerable residual risk and a functioning monitoring program is a materially different credit from one without, even at identical hazard. Its resilience, the function it keeps during an event and the speed at which it recovers, is the variable through which a physical shock reaches the balance sheets financed on it, and it is observable: recovery times after the events of Section 2 are data, not projections.

Building that bridge means working at two scales at once, because neither suffices alone. Top-down, the shocks that enter macroeconomic and financial models should be derived from physically consistent scenarios (i.e. simulated floods, droughts, heat waves and fire seasons with realistic intensity and duration) and translated into sectoral perturbations through damage functions rather than assumed. The fall in industrial output justified by the physical parameters of the damage, the loss of labor productivity derived from physiological heat indices, the tourism shock from modelled damage to coasts and access, the energy shock from modelled reductions in hydropower inflows. Bottom-up, a small number of well-documented events, studied at the local and regional scale at which their effects are visible and attributable, with damage records such as the ones collected by the Spanish Consorcio de Compensación de Seguros, sectoral data and high-resolution modelling, serve as empirical laboratories in which the full chain from hazard to household and firm to lender can be traced, and in which one can observe which channel dominates (direct damage, operational disruption, productivity loss, cascading effects through supply chains, or prices). The local cases anchor and calibrate the aggregate model; the aggregate model gives the local cases a system in which to matter. In the same loop adaptation stops being an assumption. The physical effect of a measure on the event (i.e. depth, duration, extent, water deficit) is modelled, the residual risk after adaptation is compared with the baseline, and trajectories of risk under different degrees and speeds of implementation can be placed side by side. That is what adaptive trajectories would look like in a supervisory exercise, and it is, in the author's assessment, within reach of existing methods.

The table sets out what physical-risk science can deliver to that end today, and what it cannot yet.

Table 2

What physical-risk science can and cannot yet deliver to the supervisory agenda

Can be delivered now	Cannot yet be delivered reliably
Granular hazard under multiple scenarios and horizons, including compound events and return periods in the tail.	Probabilities of emissions scenarios, and hence of the hazard projections that depend on them.
Asset-level exposure, provided assets are geolocated where they stand rather than at the registered office.	Consistent treatment of existing defenses and resilience across data providers.
Vulnerability and depth-damage functions for the main asset classes.	Calibrated firm-level adaptation and its effect on vulnerability.
Baselines of existing protection, which fix where risk stands today.	Indirect and second-order losses to territories and sectors.
Recovery times and function loss observed in past events, as the measurable resilience of a territory.	The resilience of institutions and governance, which is not a physical quantity and has to be inferred.
Adaptation pathways with tolerable residual risk and trigger conditions, where a plan exists.	The tolerable-risk decision itself, which is political and local.

One item in the left-hand column deserves emphasis. The Banco de España's own statistical work shows that allocating a firm's activity to its production centers rather than to its registered office changed the estimated impact of the DANA in one affected municipality by 17% in sales, employment and assets (Maza, 2026). Geolocation is not a detail; for physical risk it is often the difference between measuring exposure and estimating it.

6.4. The DANA as the worked example

The Valencia DANA of 29 October 2024 is the case against which the argument of this paper can be tested. It was, in the terms of Section 2, a tail event in a fat-tailed distribution, and it was locally catastrophic. By 5 June 2026 the Consorcio de Compensación de Seguros had registered 251,813 claims, 95.4% of them in the province of Valencia, had paid 212,533 of them for a cumulative €4,484 million, and estimated the total cost of the event at around €4,800 million ([CCS via SegurosNews, 2026](#)). Estimates of total economic losses have been reported at well above that figure ([Zurich, 2025](#)), which is another way of saying that the insured share was a fraction and that the rest fell on households, firms and public budgets. And yet the Banco de España assessed the event as non-systemic for the financial system, because the share of exposed credit was small and public authorities intervened (Núñez, 2026).

Every element of the bridge is present in that sentence. The transfer mechanism, a publicly administered, surcharge-funded scheme with an equalisation reserve, is itself an adaptation measure of the accommodate type, and it worked. It also separates two resiliences that the same word conceals: the territory absorbed the shock with great difficulty and was still recovering eighteen months later, when the Consorcio was still settling claims; the financial system absorbed it well. The second depends on the first only through the share of exposed credit and the speed of local recovery, this time. The public intervention that made the event non-systemic is the general-equilibrium actor doing its job. The small share of exposed credit is a partial-equilibrium fact about one event in one region. The question the DANA does not answer, and that the frequency argument of Section 2.2 forces us to ask, is what happens when the next event arrives before the reserve is rebuilt, or strikes a region where credit is more concentrated, or coincides with a macroeconomic shock. The supervisory statement of the same point, heard in the conference's closing conversation, is that climate does not create a new class of risk but drives the established ones, and that its materiality for financial institutions is amplified by the growing insurance gap. The bridge is needed not because the last event was systemic, but because the next sequence might be.

7. CONCLUSIONS: SIX MESSAGES

The question the panel was set has not been solved here. What this paper has tried to do is to structure it, and the structure yields six messages for the community the conference brought together.

1. The evidence is settled; the signal is in the extremes. Warming, ocean heat and sea-level rise are at record levels, and the rate of sea-level rise has doubled in three decades. What that implies for risk is not a shift in averages but a shift in tails, in frequency, in sequencing and in compounding (Spain's 2026 showed all four) and, through recovery time, in resilience. Analysis built on expected values will systematically understate what is coming.
2. Risk is partly a decision, and the decisions must be made explicit. Where the system boundary is drawn, whether existing protection and resilience are counted, and what residual risk is tolerated determine the number that ends up in a stress test. Users of hazard and risk data should demand to know these choices, and providers should state them.
3. Adaptation has no 1.5°C and should stop waiting for one. It is governed instead by tolerable residual risk and required resilience, how much function a system keeps and how fast it recovers, set system by system, and delivered through flexible pathways with monitored trigger points. The Belém indicators measure progress; they do not define sufficiency. Sufficiency is local.
4. The implementation gap is a governance gap. National plans and local implementation are separated by narrow funding channels and by interactions between levels of government and between public and private actors that no one manages. The European climate-resilience framework expected in late 2026 is the first opportunity to give adaptation a policy anchor that supervisors can use; it should be judged by whether it does.
5. Optimal adaptation is multi-criteria; insurability is a constraint; and the adaptation effort is itself a financial-stability question. Cost minimization and insurability cannot serve as the objective function of an adaptation plan. The scale of exposed assets exceeds what public budgets can protect, the financing structure that brings private owners, lenders and insurers into protection has not been designed, and adaptation investment compressed in time or concentrated in indebted territories generates risks of its own. A framework able to compare trajectories (i.e. inaction, disordered action, well-sequenced action) is needed and does not yet exist.
6. Adaptation is an endogenous variable in financial risk, and it can be made legible. The physical risk on a portfolio depends on public adaptation decisions the institution does not control but can observe, and on a resilience it can measure. Supervisors should ask for pathways, residual-risk statements and monitoring; scientists should deliver physically consistent scenarios, tail return periods, geolocated exposure, baselines of

existing protection and observed recovery times, calibrated on well-documented local events, and be candid about indirect losses and firm-level adaptation, which they cannot yet quantify. The right unit of analysis is the sequence of events, not the event; the right statistic is the tail, not the mean; and the right question is not whether the last event was systemic but whether the next sequence might be.

REFERENCES

AEMET. (2026a). *España registra 61 días bajo ola de calor en 2026, más que en ningún otro año*. Agencia Estatal de Meteorología, 9 September 2026. aemet.es.

AEMET. (2026b). *Un verano de cuatro meses bate récords de calor en, al menos, siglo y medio*. Agencia Estatal de Meteorología, 22 September 2026. aemet.es.

AEMET. (2026c). Informe sobre las borrascas atlánticas y las lluvias persistentes e intensas asociadas de enero y febrero de 2026. Agencia Estatal de Meteorología, 25 April 2026. aemetblog.es. Also published in *Tiempo y Clima* (AME), July 2026.

BANCO DE ESPAÑA. (2026). *Climate Risk and Sustainable Finance: Bridging Policy, Research and Practice*, High-Level Conference, Madrid, 25 March 2026. [Event page](#); [programme](#); [presentations](#); [video playlist](#).

BATTISTON, S. (2026). Remarks in Session 2, “Transition Plans and Adaptation to Physical Risk”, Banco de España High-Level Conference, 25 March 2026 (video, above).

BRUNEAU, M., CHANG, S. E., EGUCHI, R. T., LEE, G. C., O’ROURKE, T. D., REINHORN, A. M., SHINOZUKA, M., TIERNEY, K., WALLACE, W. A. and VON WINTERFELDT, D. (2003). A framework to quantitatively assess and enhance the seismic resilience of communities. *Earthquake Spectra*, 19(4), 733–752.

C3S. (2026a). *Copernicus: August was the world’s joint hottest month on record, pushing global temperatures back above 1.5°C above pre-industrial levels*. Copernicus Climate Change Service, 9 September 2026. climate.copernicus.eu.

C3S. (2026b). *Surface air temperature for August 2026*. Copernicus Climate Change Service. climate.copernicus.eu.

C3S. (2026c). *Climate Bulletins*, August 2026 edition. Copernicus Climate Change Service. climate.copernicus.eu.

CASILLAS, A. (2026). Remarks in Session 2, “Transition Plans and Adaptation to Physical Risk”, Banco de España High-Level Conference, 25 March 2026 (video, above).

CETEx. (2026). *Climate risk and sustainable finance: bridging policy, research, and practice* — event commentary, 2 April 2026. cetex.org.

CONSORCIO DE COMPENSACIÓN DE SEGUROS. (2026). Nota informativa sobre las inundaciones extraordinarias producidas por la DANA de octubre de 2024, 5 June 2026. Reported by [SegurosNews](#); see also [Gobierno de España, 26 November 2025](#) and the [CCS DANA page](#).

CONSORCIO DE COMPENSACIÓN DE SEGUROS. (2026a). Nota informativa sobre episodios de tempestad ciclónica atípica producidos entre el 20 de enero y el 20 de febrero de 2026. conorseguos.es; definitive delimitation reported by [SegurosNews, 30 June 2026](#).

EUROPEAN BANKING AUTHORITY. (2025). *Guidelines on the management of ESG risks*. EBA/GL/2025/01.

EUROPEAN COMMISSION. (2026). *European Climate Resilience and Risk Management – Integrated Framework* (forthcoming). climate.ec.europa.eu; [European Parliament Legislative Train](#).

EEA. (2024). *European Climate Risk Assessment*. EEA Report 01/2024. Copenhagen.

EFFIS. (accessed 28/10/2026). <https://forest-fire.emergency.copernicus.eu>

FERNANDEZ-PEREZ, A., LARA, J. L. LOSADA, I. J. (2025). Flexible adaptation strategies for managing compound climate risks in port infrastructures. *Coastal Engineering*, 104844.

FRANCE 24. (2026). *España vivió su verano más cálido desde que hay registros*, 9 September 2026 (EFFIS burned-area figures). france24.com.

HAASNOOT, M., KWAKKEL, J. H., WALKER, W. E. and MAAT, J. (2013). Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. *Global Environmental Change*, 23(2), 485–498.

HERSBACH, H., BELL, B., BERRISFORD, P., ET AL. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146, 1999–2049. <https://doi.org/10.1002/qj.3803>

IPCC. (2019). *Special Report on the Ocean and Cryosphere in a Changing Climate*. Cambridge University Press.

IPCC. (2021). *Climate Change 2021: The Physical Science Basis. Summary for Policymakers*. Working Group I contribution to the Sixth Assessment Report. Cambridge University Press.

IPCC. (2022). *Climate Change 2022: Impacts, Adaptation and Vulnerability*. Working Group II contribution to the Sixth Assessment Report; Cross-Chapter Paper 4, Mediterranean Region. Cambridge University Press.

LOSADA, I. J., FELIU, E. and SANZ, M. J. (coords.) (2025). *Evaluación de Riesgos e Impactos derivados del Cambio Climático en España (ERICC-2025)*. Ministerio para la Transición Ecológica y el Reto Demográfico, Madrid. [Executive summary](#); [sectoral chapter on the financial system and insurance](#)

LOSADA, I. J., TOIMIL, A., MUÑOZ, A., GARCIA-FLETCHER, A. P., DIAZ-SIMAL, P. (2019). A planning strategy for the adaptation of coastal areas to climate change: The Spanish Case. *Ocean & Coastal Management*, 182, 104983.

MAGNAN, A. K., BELL, R., DUVAT, V. K. E, FORD, J. D., GARSCHAGEN, M., HAASNoot, M., LACAMBRA, C., LOSADA, I. J. (and 10 other authors). (2023). Status of global coastal adaptation. *Nature Climate Change*, 13, 1213-1221.

MARKANDYA, A. (2026). *Climate Risk and Sustainable Finance*. Keynote presentation, Banco de España, 25 March 2026 (in the conference presentations pack, above; figures drawn from Climate Policy Initiative data).

MAZA, L. Á. (2026). *Integrating the climate dimension into economic statistics: the experience of the Banco de España*. Presentation, Session 1 (in the presentations pack, above).

MITECO. (2017). *Estrategia de Adaptación al Cambio Climático de la Costa Española*. Dirección General de Sostenibilidad de la Costa y del Mar, approved 24 July 2017. miteco.gob.es.

NGFS. (2025). *NGFS short-term climate scenarios*. Network for Greening the Financial System.

NÚÑEZ, S. (2026). *Climate risk and sustainable finance: bridging policy, research and practice*. Opening remarks, Banco de España, 25 March 2026. [BIS central bankers' speeches](#).

UNEP. (2025). *Adaptation Gap Report 2025: Running on Empty*. Nairobi. ([summary](#)).

UNFCCC. (2025). Decisions of COP30/CMA7, Belém, November 2025 (Global Goal on Adaptation indicators; adaptation finance). See [Carbon Brief](#) and [Adaptation Community](#).

WMO. (2025). *State of the Global Climate 2024*. WMO-No. 1368. Geneva. ([summary of findings](#))

WMO. (2026). *State of the Global Climate 2025*. WMO-No. 1391. Geneva, 23 March 2026. [wmo.int](#).

ZURICH INSURANCE. (2025). *Valencia floods: lessons in claims and recovery, one year on*. [zurich.com](#).