H HEURISTICS

Navigating a Changing World

Building Resilient Development Pathways Through the Global Polycrisis

How Adaptation and Clean Infrastructure Can Reduce Systemic Vulnerability

Why the Pathway a Developing Economy Chooses Now Locks In Decades of Vulnerability — or Resilience

H Heuristics · Hunter Hughes

Research Report · September 2026

Abstract

The defining danger of the current period is not the size of any single hazard but the coupling between hazards, and that coupling falls hardest on the economies least able to absorb it. Since 2020 a pandemic, a food and energy price spike, the sharpest rise in global borrowing costs in a generation, an accelerating run of climate extremes, and a fragmenting geopolitics have arrived on top of one another and transmitted their stress into one another. This report treats that condition, commonly labelled "polycrisis", as a development problem: shocks that originate in one system now propagate through a small set of shared channels — prices, exchange rates, fiscal space, credit, and confidence — that are concentrated in the structure of developing economies, so that the same physical hazard does far more systemic damage there than in a diversified, local-currency economy.

The report's central argument is that the systemic vulnerability which turns coupled shocks into cascading crises is built rather than given — accumulated through the composition of a country's energy system, the placement and standard of its infrastructure, the depth of its safety nets, and the state of its public finances — and can therefore be reduced by deliberate choices about how development proceeds. Using the disaster-risk framework of hazard, exposure, sensitivity, and adaptive capacity, it identifies two levers that do most of that work: adaptation, which raises adaptive capacity and lowers sensitivity to the shocks that arrive regardless; and clean infrastructure, which cuts exposure, severs the imported-fuel transmission channel, hardens systems through distribution, and shrinks the hazard itself. The two are complements whose combined return exceeds the sum of their parts, concentrated in a set of double-duty investments that pay in conventional development terms and in resilience at once.

The binding constraints are financial and institutional, not technical. The highest-return interventions are known and their benefit–cost ratios are exceptional: the Global Commission on Adaptation put returns across a portfolio of measures at two-to-one up to ten-to-one, the World Bank's Lifelines analysis found roughly four dollars of benefit per dollar of resilient-infrastructure premium, and multi-hazard early warning returns on the order of ten to one. Yet UNCTAD reports developing-country external debt at a record US$11.4 trillion and net interest at US$921 billion in 2024, with 3.4 billion people in countries that spend more on debt service than on health or education; the UNEP Adaptation Gap Report 2025 puts adaptation needs at US$310–365 billion a year by 2035 against international public flows of about US$26 billion; and clean-energy investment, now roughly twice that in fossil fuels globally, barely reaches developing economies outside China because of a 400–800 basis-point cost-of-capital wedge that a capital-intensive, front-loaded pathway is acutely sensitive to.

Because infrastructure is long-lived and vulnerability accumulates, the timing of the choice is decisive: the pathway a developing economy sets in this decade is locked in for the life of the assets it builds, and the premium on the resilient option is smallest before the concrete is poured. The report sets out the framework, shows why the polycrisis is a development problem, examines lock-in and each of the two levers, makes the case for their synergy, addresses the financing and institutional constraints, proposes a three-tier sequencing rule, and grounds the argument in the experience of Bangladesh, Kenya, Morocco, Vietnam, Barbados, and Rwanda. Its central finding is that the countries which absorb shocks well are distinguished less by their wealth or exposure than by whether they built and protected the enabling capacity — fiscal space, delivery institutions, pre-arranged finance, durable coordination — before the shocks arrived; and that the task for the coming decade is not to invent a resilience agenda but to fund it, integrate it, sequence it, and build it in advance.

Keywords: polycrisis; systemic vulnerability; risk transmission; development pathways; path dependence; infrastructure lock-in; climate adaptation; adaptation finance; clean infrastructure; distributed energy; energy

security; cost of capital; sovereign debt; adaptive social protection; early warning systems; nature-based solutions; anticipatory action; maladaptation; benefit–cost ratio

JEL classification: Q54, O13, H54, F34, O19; climate and natural disasters; agriculture, natural resources and environment in development; infrastructure and public investment; international lending and sovereign debt; economic development and international linkages

Data and method: This report synthesises institutional and peer-reviewed research on systemic risk, climate adaptation, clean infrastructure, and development finance, including the *IPCC Sixth Assessment Report* (WGII); the WMO *Atlas of Mortality and Economic Losses (1970–2021)* and the *UNDRR Global Assessment Report*; the Global Commission on Adaptation's *Adapt Now* and the World Bank's *Lifelines*; the *UNEP Adaptation Gap Report 2025*; UNCTAD's *A World of Debt 2025*; the *IEA World Energy Investment 2025* and *SDG7* data; the World Bank *State of Social Protection Report 2025*; and primary reporting on country programmes in Bangladesh, Kenya, Morocco, Vietnam, Barbados, and Rwanda. Every quantitative claim is attributed inline to a primary or authoritative secondary source. Figures 1, 5, and 6 are conceptual schematics; Figure 2 is an illustrative decadal series reflecting documented WMO trends; Figures 4 and 8 present indicative ranges. The report is analytical rather than predictive.

Contents

1. Introduction: Development at Risk in an Age of Polycrisis

> *For most of the postwar era, development was a story about getting on to a path. The argument of this report is that the harder question now is which path — because the pathways open to a developing economy differ not only in how fast they raise incomes, but in how much systemic vulnerability they build in along the way.*

The first quarter of the twenty-first century has delivered an unusually dense run of shocks. A pandemic became, in short order, an education crisis, a debt crisis, and a supply-chain crisis. A war in Europe became a food crisis and an energy crisis that raised prices in countries thousands of kilometres from the front line. A cycle of interest-rate rises intended to tame inflation in the advanced economies became, across dozens of developing countries, a fiscal crisis that is still crowding out spending on the very things — health, education, adaptation — that make the next shock survivable. And running underneath all of it, a steadily rising tide of climate extremes has turned what were once tail risks into recurring line items in national budgets.

The word that has attached itself to this pattern is "polycrisis": the simultaneous, interacting operation of crises across distinct systems, such that the combined effect exceeds the sum of the parts. The term has become common enough to be worn smooth, and it is sometimes used to mean little more than "many bad things at once." This report uses it more precisely. The claim is not that the world faces an unusual number of problems; it is that the problems are increasingly *coupled* — a shock in one system now propagates rapidly into others through a small set of shared channels: prices, exchange rates, fiscal space, credit, and confidence. The coupling, not the individual shocks, is what makes the current period dangerous.

That distinction matters most for developing economies, because they sit where the channels converge. They import food and fuel priced in dollars, borrow in dollars at a premium, depend on a narrow base of commodity or tourism revenue, carry thin fiscal buffers, and are disproportionately exposed to the physical hazards that climate change is amplifying. When shocks arrive together, these economies have the least room to absorb any one of them and the fewest instruments to stop one from becoming several. The result is visible in the data: after decades of convergence, the poorest countries have, since 2020, been falling behind again — what the World Bank has called a period of reversal for the economies that can least afford it.

But vulnerability is not destiny, and it is not fixed. This report treats systemic vulnerability as something that is *built* — accumulated through the composition of a country's energy system, the placement and standard of its infrastructure, the depth of its safety nets, and the state of its public finances — and therefore something that can be *un-built* by deliberate choices about how development proceeds. The central argument is that two levers do most of that work. The first is **adaptation**: the set of investments and institutions that reduce how much damage a given shock does and how quickly a country recovers. The second is **clean infrastructure**: the reconstruction of energy and physical systems in a form that is at once lower-carbon, cheaper to run, more distributed, and harder to knock over.

Neither lever is new, and neither is proposed here as a silver bullet. What is new is the case for treating them together, and for treating them as core development strategy rather than as environmental add-ons paid for out of whatever is left over. The report will argue that adaptation and clean infrastructure are unusual among policy options because they pay twice — once in conventional development terms (growth,

access, productivity, health) and again in reduced exposure to the coupled shocks that now define the risk environment. Investments that pay twice are exactly what a capital- and capacity-constrained government should prioritise, and they are the reason a resilient pathway need not be a slower one.

The report proceeds as follows. Section 2 sets out a working framework for systemic vulnerability, distinguishing hazard, exposure, sensitivity, and adaptive capacity, and locating where the two levers act. Section 3 shows why the polycrisis is, in its incidence, a development problem. Section 4 examines lock-in — the reason the pathway chosen in this decade persists for many decades after. Sections 5 and 6 take the two levers in turn; Section 7 makes the case that they reinforce each other and identifies the double-duty investments where their returns overlap. Sections 8 and 9 address the two binding constraints, finance and institutions. Section 10 proposes a practical sequencing rule, Section 11 grounds the argument in country experience, and Section 12 turns to measurement and the real danger of maladaptation. Section 13 concludes.

2. From Hazard to Systemic Vulnerability: A Working Framework

> *Before asking how to reduce vulnerability, it helps to be precise about what vulnerability is. Loose usage treats it as a synonym for danger; the more useful definition, drawn from disaster-risk science, breaks it into parts that can be measured and acted on separately.*

The framework this report uses follows the risk concept adopted by the Intergovernmental Panel on Climate Change: risk arises from the interaction of three elements — a **hazard**, the **exposure** of people and assets to that hazard, and their **vulnerability** to harm when exposed. Vulnerability in turn has two components: **sensitivity**, how much damage a given shock inflicts on an exposed system, and the inverse of **adaptive capacity**, the ability to anticipate, cope with, absorb, and recover from that damage. A drought is a hazard; the farmers in its path are the exposure; whether their crops fail, and whether a failed harvest becomes a hunger crisis or merely a hard year, is a matter of sensitivity and adaptive capacity.

The value of the decomposition is that it separates the part of risk we cannot easily change from the parts we can. In the short run, a small developing economy has essentially no control over the global hazard climate — over whether the next monsoon is erratic or the next El Niño severe. What it can influence, and influence considerably, is how many people and how much productive capital sit in harm's way, how much damage a shock does when it lands, and how fast the system recovers afterwards. Almost every effective resilience policy works on one of these three: it moves exposure, lowers sensitivity, or raises adaptive capacity. Figure 1 places the two levers of this report against that anatomy.

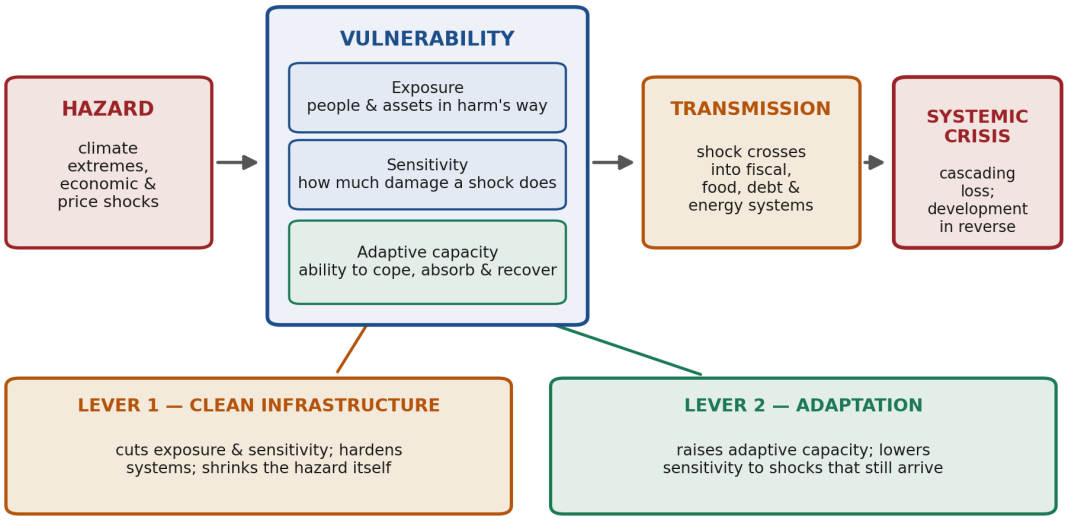

The two levers reduce systemic vulnerability by acting on its determinants — before the next shock arrives.

Figure 1. *The anatomy of systemic vulnerability, and where the two levers act.*
Framework after the IPCC risk concept (hazard × exposure × vulnerability). Schematic by the author.

2.1 What makes vulnerability *systemic*

The adjective "systemic" adds a fourth idea to the standard framework: transmission. A shock is systemically dangerous not merely because it is large where it lands, but because it can cross from the system it strikes into others that appeared unrelated. A flood that destroys a port is a logistics problem; when it interrupts fuel imports, spikes food prices, forces emergency spending that widens the deficit, triggers a credit downgrade that raises borrowing costs, and thereby shrinks the fiscal space available for the next flood, it has become a systemic problem. The damage is done less by the water than by the couplings — the fixed channels along which stress travels between systems.

This is why vulnerability cannot be assessed sector by sector. A power system, a debt profile, a food-import bill, and a social safety net can each look adequate in isolation and still combine into fragility, because they are linked: the safety net depends on the fiscal space, the fiscal space depends on the import bill, the import bill depends on the energy system. Reducing systemic vulnerability therefore means attending both to the strength of individual systems and to the couplings between them — either loosening the couplings so a shock cannot travel, or building enough capacity in each system that an arriving shock is absorbed rather than passed on.

The two levers map cleanly onto this task. Clean infrastructure works mainly on the left of Figure 1: it reduces exposure and sensitivity — by hardening assets, decentralising them, and removing the imported-fuel dependency that is one of the most reliable transmission channels — and it also shrinks the hazard itself, since the emissions avoided today lower the climate hazard everyone faces tomorrow. Adaptation works mainly on the right: it raises adaptive capacity and lowers sensitivity for the shocks that arrive regardless, through early warning, resilient design, pre-arranged finance, and safety nets that expand on cue. Together they attack vulnerability from both directions.

2.2 A note on what the framework is not

Two clarifications guard against misreading. First, the framework is not a claim that hazard reduction — mitigation — can be neglected. Every tonne of emissions avoided lowers the future hazard, and for that reason mitigation and adaptation are not rivals but complements; clean infrastructure is valuable precisely because it does both. Second, the framework is descriptive, not deterministic. Vulnerability is a *predisposition* to harm, not a prediction of it. Two countries with identical exposure can experience the same hazard very differently depending on capacities built years earlier — which is the entire practical point. The chapters that follow are, in effect, an argument about how to build those capacities before they are needed.

3. The Polycrisis Is a Development Problem

> *The polycrisis is often described in global terms, as if it fell evenly on the world. It does not. Its incidence is sharply skewed toward the countries with the least capacity to absorb it, which makes it, above all, a development problem.*

Start with the hazard trend itself. The World Meteorological Organization, reviewing half a century of records, found that the number of reported weather, climate, and water-related disasters rose roughly fivefold between the 1970s and the 2010s, and that reported economic losses rose roughly sevenfold, to a cumulative US$4.3 trillion over 1970–2021. The one genuinely encouraging line in that record is mortality, which fell markedly over the same period — overwhelmingly because of improved early warning and disaster preparedness. That divergence, more hazards and larger losses but fewer deaths, is the single most important empirical fact in this report: it shows that vulnerability is malleable, and that the instruments which bend it are already known. Figure 2 shows the pattern.

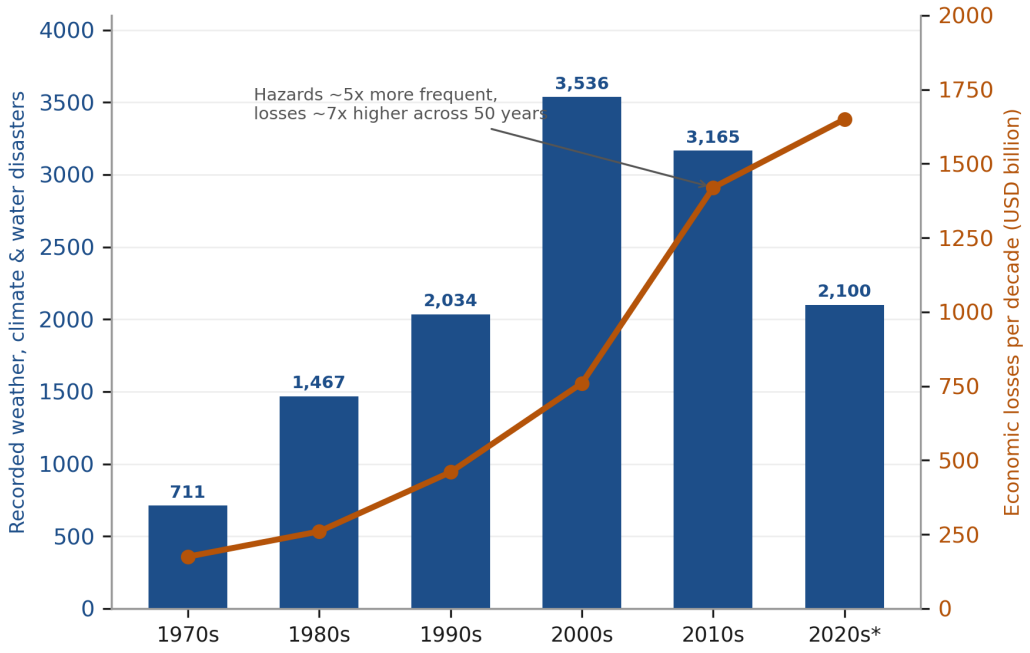

Figure 2. *Recorded weather, climate, and water disasters and economic losses per decade.*
Illustrative decadal series reflecting documented trends in the WMO Atlas of Mortality and Economic Losses (1970–2021). Values approximate; 2020s partial.

The losses, however, do not fall evenly. In absolute terms the largest dollar losses are recorded in rich countries with more to destroy; but as a *share of the economy*, and in loss of life and lasting damage to development, the burden falls overwhelmingly on low- and middle-income countries. A single tropical cyclone can erase a decade of growth in a small island state; a multi-season drought can push a low-income economy from slow progress into contraction. Because these countries have thinner buffers, the same physical shock does more relative damage — higher sensitivity — and because they have less capacity to respond, recovery is slower — lower adaptive capacity. The hazard is global; the vulnerability is developmental.

3.1 Where the channels converge

The transmission channels described in Section 2 are not evenly distributed either; they are concentrated in the structural features of developing economies. Four features stand out. First, **import dependence** in food and fuel priced in dollars means a global price spike lands directly on the household budget and the trade balance at once. Second, **commodity or tourism concentration** on the revenue side means a single external shock can hollow out government income. Third, **hard-currency debt** means a rise in global interest rates or a currency depreciation raises the real cost of the existing debt stock without any new borrowing. Fourth, **thin fiscal and institutional buffers** mean there is little slack to absorb any of the above. Each feature is a channel; together they are why a shock that would be a manageable disturbance in a diversified, local-currency economy becomes a systemic event in a developing one.

The debt channel deserves particular emphasis because it has tightened dramatically. UNCTAD reports that developing-country external debt reached a record US$11.4 trillion, and that net interest payments hit US$921 billion in 2024. Most striking, it estimates that 3.4 billion people now live in countries that spend more on debt-service interest than on either health or education. When so much of a budget is pre-committed to creditors, the fiscal space that would otherwise absorb a shock — or fund the resilience that prevents one — simply is not there. Debt is not a separate crisis sitting alongside the climate and energy crises; it is the channel through which those crises become unrecoverable.

3.2 Development in reverse

The consequence is visible in the development record. The long post-1990 trend of poorer countries slowly catching up to richer ones stalled after 2020, and for the poorest group it reversed. Progress on extreme poverty, which had fallen for a generation, flattened; a number of low-income countries emerged from the pandemic-and-price-shock years with lower real incomes, higher debt, and depleted reserves than they entered with. This is what "development in reverse" means concretely: not a single dramatic collapse, but the quiet erosion of hard-won gains as each shock consumes the resources that would have funded the next stage of progress and the defences against the next shock.

It is precisely because the losses compound in this way that the framing of this report matters. If the shocks were one-off and uncorrelated, the right response would be insurance and relief — absorb the hit and rebuild. But when shocks are frequent, coupled, and rising, the arithmetic of reactive response fails: a country cannot borrow its way through an indefinite series of emergencies at a 400-to-800 basis-point premium. The only durable answer is to lower the vulnerability that turns each hazard into a systemic loss — which returns us to the two levers, and to the reason their timing is so unforgiving.

4. Lock-In: Why the Pathway Chosen Now Persists for Decades

> *If vulnerability could be adjusted freely at any moment, its current level would matter less. The reason the coming decade is decisive is that the choices made in it are, quite literally, built to last — and unbuilding them is far more expensive than building them differently the first time.*

Development is a physical process. It lays down power plants, transmission networks, roads, ports, water systems, and — through zoning, land tenure, and the sheer momentum of where people settle — the shape of cities themselves. These assets are extraordinarily long-lived. A coal or gas plant runs for around forty years; major water and irrigation systems for decades; roads, ports, and grids for half a century or more; the built form and land-use pattern of a city can persist, in its essentials, for a century. Figure 3 sets these lifetimes against the narrow window in which the defining choices are actually made.

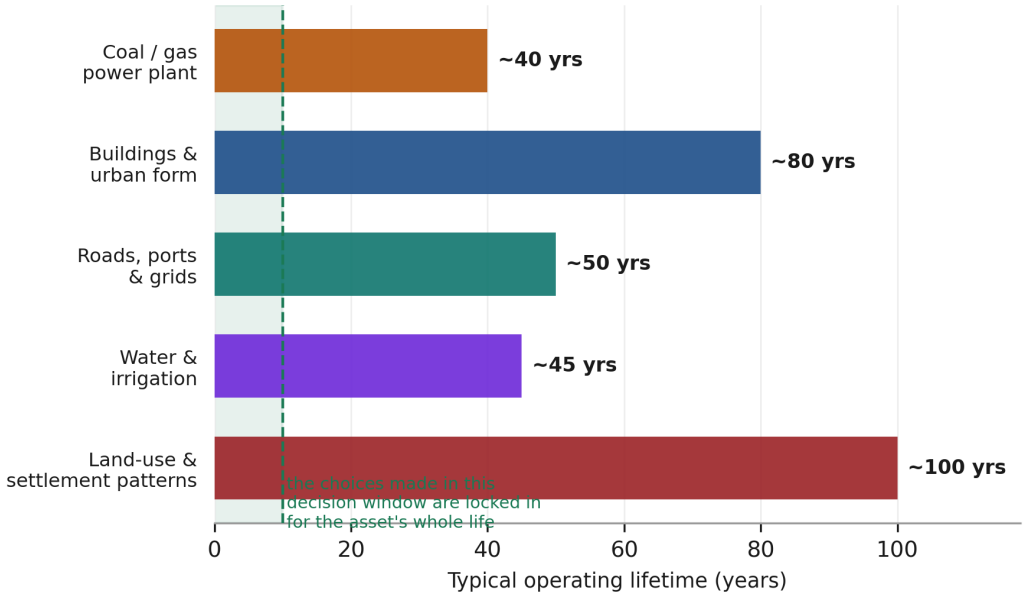

Figure 3. *Typical operating lifetimes of core infrastructure, against the near-term decision window.*
Indicative asset lifetimes compiled by the author from standard engineering and planning ranges.

The implication is a form of path dependence economists call lock-in. When a country builds a fossil-fuelled, centralised, hazard-exposed infrastructure base, it does not merely emit for the life of those assets; it also commits itself to their vulnerabilities — to the imported-fuel price exposure, the single points of failure, the floodplain siting — for decades, because ripping them out early means writing off capital a poor country cannot afford to waste. This is the trap of the "grow now, clean up later" pathway: the cleanup, when it finally comes, must be paid for twice, once to build the vulnerable asset and again to replace it before its time.

4.1 The stranded-asset problem in reverse

Rich countries confront lock-in as a stranded-asset problem: they must retire fossil capital before the end of its economic life to meet climate goals, and absorb the loss. Many developing countries face a different and, in one respect, more fortunate version of the same problem. Because much of the infrastructure they

will use in 2050 has not yet been built, they are not primarily unwinding bad past choices; they are deciding, now, whether to make them. The share of a developing economy's mid-century capital stock that is still on the drawing board is far larger than in a mature economy — which means the marginal cost of choosing the resilient option is far lower, because it is the difference between two new builds rather than the full cost of an early demolition plus a replacement.

This is the optimistic core of the lock-in argument, and it is easy to miss beneath the pessimism about how little time remains. The window is short, but it is open, and it is open widest exactly where the future asset base is largest relative to the existing one. A country electrifying for the first time can build distributed and clean rather than centralised and fossil at little or no cost penalty and a substantial resilience gain. A city still taking shape can steer settlement away from the floodplain far more cheaply than one that must later defend or relocate it. The premium on resilience is smallest before the concrete is poured.

4.2 Why markets under-supply resilience

If the resilient pathway is so advantageous, why is it not simply chosen? Part of the answer is finance, addressed in Section 8. But part is structural: markets systematically under-supply resilience because its benefits are diffuse, deferred, and probabilistic while its costs are concentrated, immediate, and certain. A resilient bridge costs more today for the benefit of surviving a flood that may or may not come in twenty years, and whose avoided cost accrues to the public rather than the builder. Much of the value of resilience is an avoided loss that, if the investment works, is never observed — a notoriously hard thing to justify in a budget cycle or to a lender pricing a short-tenor loan. Left to itself, the market builds the cheap, vulnerable version. Correcting that bias — through standards, public co-investment, and the pricing of avoided losses — is a core task of resilience policy, and a recurring theme in the sections that follow.

5. The First Lever: Adaptation as Vulnerability Reduction

> *Adaptation has a public-relations problem: it sounds like defeat, an admission that mitigation has failed. Framed correctly it is nothing of the kind. It is the highest-return category of development spending available to a shock-exposed economy — and the evidence for that claim is unusually strong.*

Adaptation is the deliberate reduction of sensitivity and the deliberate raising of adaptive capacity: everything a society does to ensure that the hazards it cannot prevent do less damage and are recovered from faster. It ranges from the physical (sea walls, drainage, climate-resilient roads and crop varieties) to the informational (early-warning systems, hydrological monitoring) to the financial and institutional (pre-arranged disaster finance, insurance, safety nets that scale up on cue). What unites these otherwise disparate measures is that each intervenes on the right-hand side of Figure 1, where vulnerability is most tractable.

5.1 The returns are exceptional and well-documented

The economic case for adaptation is not speculative. The Global Commission on Adaptation, in its landmark *Adapt Now* assessment, estimated that investing US$1.8 trillion across five areas — early-warning systems, climate-resilient infrastructure, improved dryland agriculture, mangrove protection, and more resilient water resources — over 2020–2030 could generate US$7.1 trillion in total net benefits, with benefit–cost ratios ranging from roughly two to one up to ten to one depending on the measure. The World Bank's *Lifelines* report reached a complementary conclusion for infrastructure specifically: investing in more resilient infrastructure in low- and middle-income countries would yield some US$4.2 trillion in net benefits — about four dollars back for every dollar of the resilience premium.

Early warning is the standout. Multi-hazard early-warning systems are among the highest-return investments in all of public policy, with benefit–cost ratios commonly estimated in the range of nine or ten to one, because a modest, cheap, standing capability averts enormous, expensive losses. The mechanism is simple: twenty-four hours' warning of a storm or heatwave allows people and movable assets to get out of the way, cutting damage by roughly a third. This is the instrument most responsible for the falling mortality in Figure 2, and it is the flagship of the UN's Early Warnings for All initiative, which aims for universal coverage by 2027 — a target that remains, as of the mid-2020s, only about half met. Figure 3's message about timing applies here too: the cheapest life to save is the one saved by a warning system built before the storm.

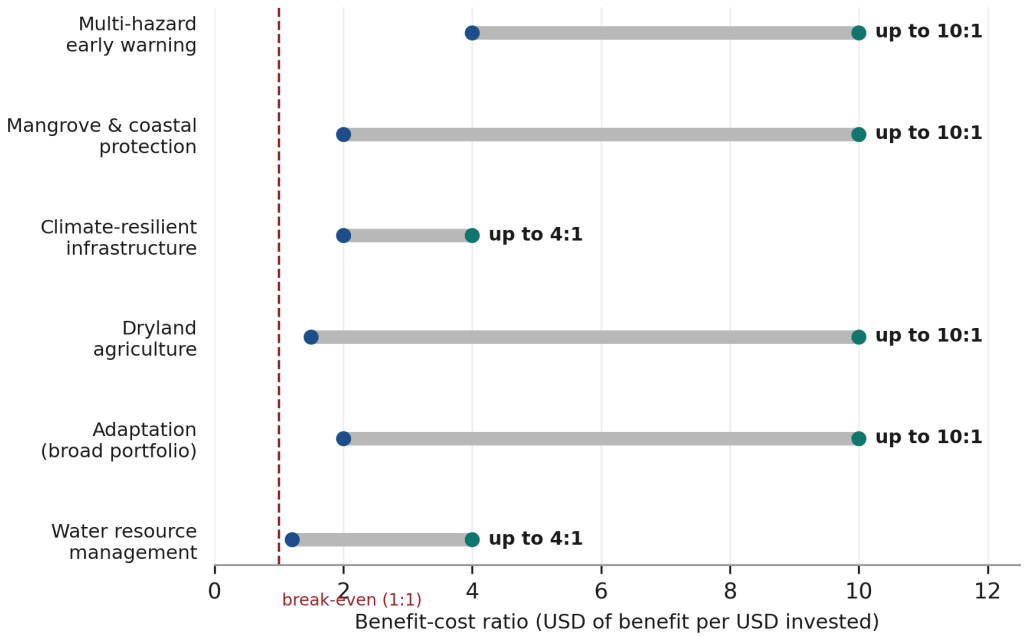

Figure 4. *Benefit–cost ratios of major resilience investments (ranges).*

Ranges drawn from the Global Commission on Adaptation (Adapt Now, 2019) and the World Bank (Lifelines, 2019). Values indicative.

Table 1. What works in adaptation: representative interventions and how they reduce vulnerability.

Intervention	Determinant it acts on	How it reduces systemic vulnerability	Indicative BCR
Multi-hazard early warning	Sensitivity; adaptive capacity	Moves people and movable assets out of harm's way; averts mortality and damage	up to ~10:1
Climate-resilient infrastructure	Sensitivity	Keeps roads, power, and water functioning through shocks; prevents service cascades	~2–4:1
Adaptive social protection	Adaptive capacity	Delivers cash on cue so a shock does not force asset sales or hunger	high in shock years
Nature-based solutions (mangroves, wetlands)	Exposure; sensitivity	Absorbs storm surge and flood; also stores carbon and supports livelihoods	up to ~10:1
Dryland & climate-resilient agriculture	Sensitivity	Stabilises yields and food supply through drought and heat	up to ~10:1
Pre-arranged disaster finance	Adaptive capacity	Funds response and recovery immediately, avoiding costly emergency borrowing	cuts response cost sharply

Benefit–cost ratios are indicative ranges from the sources cited in the text; actual returns are context-specific.

5.2 The most productive adaptation is often institutional, not physical

It is tempting to picture adaptation as concrete — sea walls and raised roads. But some of the highest-leverage measures are institutional and financial, because they raise adaptive capacity across every hazard at once. **Adaptive social protection** is the clearest example: a safety-net system with a registry of households, a digital payment rail, and a pre-agreed trigger can expand automatically when a drought index crosses a threshold, delivering cash to affected families within days rather than months. That speed is

decisive, because it is the interval between a shock and a response that determines whether a household copes by drawing on savings or by selling its productive assets, pulling children from school, and cutting meals — the coping strategies that turn a bad season into a permanent setback.

Pre-arranged finance works on the same logic at the level of the state. A government that has secured contingent credit, disaster-risk insurance, or a debt instrument that pauses repayment after a catastrophe can respond immediately and at known cost, rather than scrambling to borrow at emergency rates while a downgrade raises the cost of everything else. The instrument matters less than the timing: money arranged before the shock is worth several times the same money arranged after it. This is why the financial architecture discussed in Section 8 is properly understood as adaptation infrastructure, every bit as much as a drainage canal.

5.3 The gap between what pays and what is funded

Given returns like these, the striking fact is how little adaptation is actually financed. The UN Environment Programme's *Adaptation Gap Report* estimates that developing countries will need on the order of US$310–365 billion a year for adaptation by 2035, against international public adaptation-finance flows of roughly US$26 billion — a gap of twelve to fourteen times. The shortfall is not explained by ignorance of what works or by poor returns; both are well established. It is explained by the financing constraint of Section 8 and the institutional constraint of Section 9. Adaptation is the rare case where the analysis is settled and the bottleneck is entirely in delivery.

6. The Second Lever: Clean Infrastructure as Resilience

> *The case for clean infrastructure is usually made on climate grounds. That case is sound, but it undersells the point. In a developing economy exposed to coupled shocks, clean infrastructure is also, and often primarily, a resilience investment — one that happens to cut emissions as a by-product.*

"Clean infrastructure" here means more than renewable power plants. It is the reconstruction of a country's energy and physical systems in a form that is lower-carbon, cheaper to operate, more distributed, and more robust: solar and wind generation, storage, modern grids and mini-grids, electrified transport and cooking, and the efficient, well-sited buildings and networks that go with them. Its resilience value comes from four properties, each of which acts on a transmission channel identified in Section 3.

6.1 Four ways clean infrastructure reduces vulnerability

- **It severs the imported-fuel channel.** A country that generates power from domestic sun and wind is not exposed, on that portion of its energy, to the dollar-denominated fuel-price spikes that transmit a distant war into a domestic fiscal and balance-of-payments crisis. Solar and wind have no fuel bill and therefore no fuel-price risk — a direct reduction in one of the most reliable couplings in the system.
- **It is inherently more distributed.** Rooftop solar, community mini-grids, and storage spread generation across thousands of nodes rather than concentrating it in a few large plants and long transmission lines. A distributed system has no single point of failure a storm or attack can take out; when part fails, the rest keeps running. Centralisation is a sensitivity multiplier, and distribution is its antidote.
- **It lowers operating costs permanently.** Because the marginal fuel cost of solar and wind is zero, clean generation converts a recurring, volatile operating expense into a fixed, up-front capital cost. For a government, that predictability is itself resilience: a shock cannot inflate a bill that has no fuel in it.
- **It shrinks the hazard.** Uniquely among the interventions in this report, clean infrastructure works on the far left of Figure 1 as well: every tonne of emissions it avoids lowers the future climate hazard that drives the physical shocks everyone else must adapt to. It is simultaneously mitigation and adaptation — which is exactly why the mitigation-versus-adaptation framing is a false choice.

The economics that make this possible have shifted decisively. Solar photovoltaics and onshore wind are now the cheapest source of new electricity in most of the world, and the International Energy Agency reports that global clean-energy investment reached roughly US$2.2 trillion in 2025 — about twice the roughly US$1.1 trillion going to fossil fuels — with solar PV alone the single largest category at around US$450 billion. The technology is no longer the constraint. For a country building its energy system for the first time, the clean, distributed option is now frequently the cheapest option on a levelised basis, before its resilience premium is even counted.

6.2 The catch: clean capital barely reaches the countries that need it

The aggregate figures conceal a distribution problem that is central to this report. That US$2.2 trillion is overwhelmingly concentrated in the advanced economies and China; the share flowing to developing

economies outside China is small and has been slow to grow, even though those countries hold much of the world's best solar and wind resource and most of its unmet energy demand. The IEA and others attribute this not to a shortage of viable projects but to the cost of capital: the same clean-energy project can face borrowing costs several times higher in a developing economy than in an advanced one, a wedge commonly put at 400 to 800 basis points. Because clean infrastructure is capital-intensive and front-loaded — its cost is almost all up-front finance and almost none is fuel — it is exceptionally sensitive to the cost of that finance. A high cost of capital does more to keep the clean pathway out of reach than any technological limit.

This is the hinge on which the two levers meet the financing constraint. The resilient pathway is not blocked by physics or even, any longer, by the price of the technology. It is blocked by the price of money — and, as Section 8 argues, that price is substantially a function of choices about how capital is structured and de-risked, which means it is a policy variable, not a fact of nature.

6.3 Access as the first resilience

For the roughly 660 million people who still lack any electricity access and the some two billion without clean cooking, the resilience argument begins even earlier. Energy access is itself adaptive capacity: refrigeration for medicine and food, power for pumping and irrigation, connectivity for early warning, cooling in rising heat. Programmes such as the World Bank and African Development Bank's Mission 300, which aims to connect 300 million Africans to electricity by 2030, are therefore resilience programmes as much as development ones — and because so much of that new access will be built from scratch, it is the clearest case of the lock-in opportunity in Section 4. The households being connected now will be connected, cleanly or otherwise, for decades; the marginal cost of doing it in the resilient form is at its lowest exactly at the moment of first build.

7. Why the Two Levers Reinforce Each Other

> *The strongest argument for the two levers is not that each is individually worthwhile — Sections 5 and 6 establish that — but that together they are worth more than the sum of their parts. Pursued as one agenda, adaptation and clean infrastructure compound; pursued separately, they leave value on the table.*

Recall the anatomy of Section 2. Clean infrastructure works mainly on exposure, sensitivity, and the hazard itself; adaptation works mainly on adaptive capacity and residual sensitivity. Because they act on different determinants of the same quantity, their effects multiply rather than merely add. A clean, distributed power system reduces how often and how badly the grid fails; adaptive institutions ensure that when it does fail, service is restored and affected households are supported quickly. Neither alone delivers a resilient energy system; together they do. The same complementarity holds across the board: resilient roads (adaptation) are worth more when they carry electrified transport freed from fuel-price shocks (clean infrastructure); early-warning systems (adaptation) run on the reliable power and connectivity that clean, distributed energy provides.

The combined effect is what Figure 5 depicts. On a business-as-usual pathway, systemic vulnerability rises as hazards intensify and lock-in accumulates. On a resilient pathway, the two levers bend the curve downward over time — first flattening the rise as clean infrastructure cuts exposure, then pulling it down as adaptive capacity deepens. The shaded area between the curves is the avoided vulnerability: the cascades that never happen, the development gains that are not reversed, the recoveries that are faster and cheaper. It is also, in a real sense, the return on the strategy.

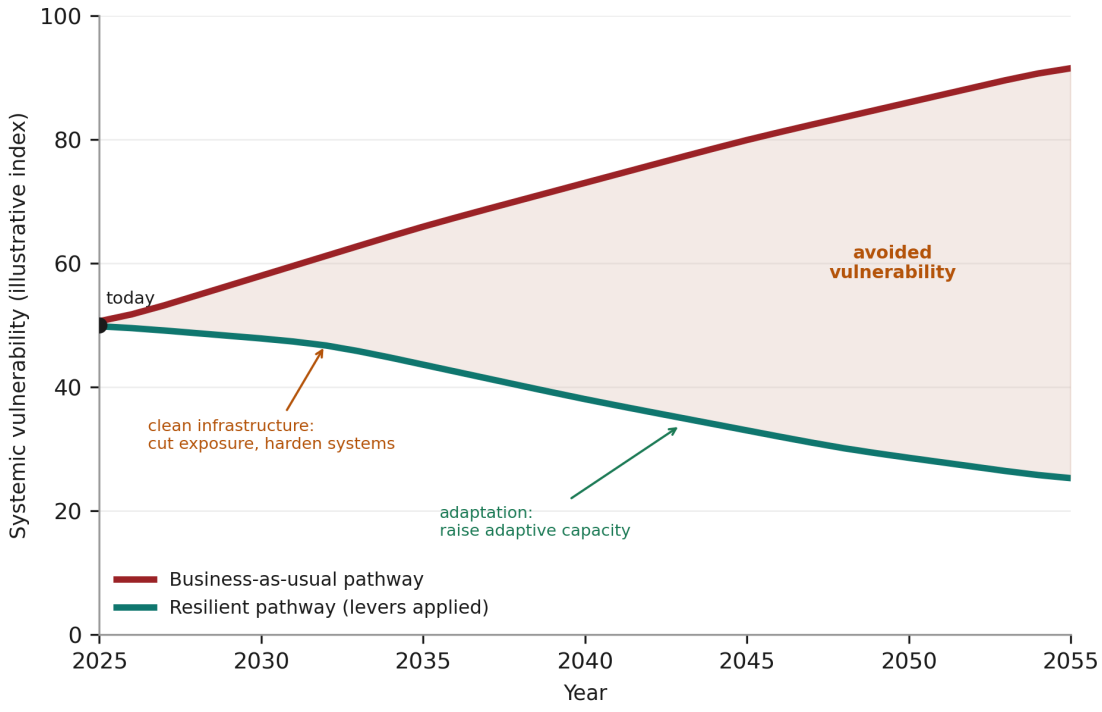

Figure 5. *Two development pathways: systemic vulnerability under business-as-usual and under the two levers.*
Illustrative schematic by the author; the vulnerability index is conceptual, not an empirical projection.

7.1 Double-duty investments: where the returns overlap

The practical significance of the synergy is that certain investments strengthen both levers at once — and it is on these that a capital- and capacity-constrained government should concentrate. Call them double-duty investments: they deliver a conventional development payoff (growth, access, productivity, health) *and* a vulnerability reduction, so that they can be justified on development grounds even by a government that does not price resilience at all, while quietly buying resilience as well. Figure 6 maps a set of candidates against both axes; Table 2 details them.

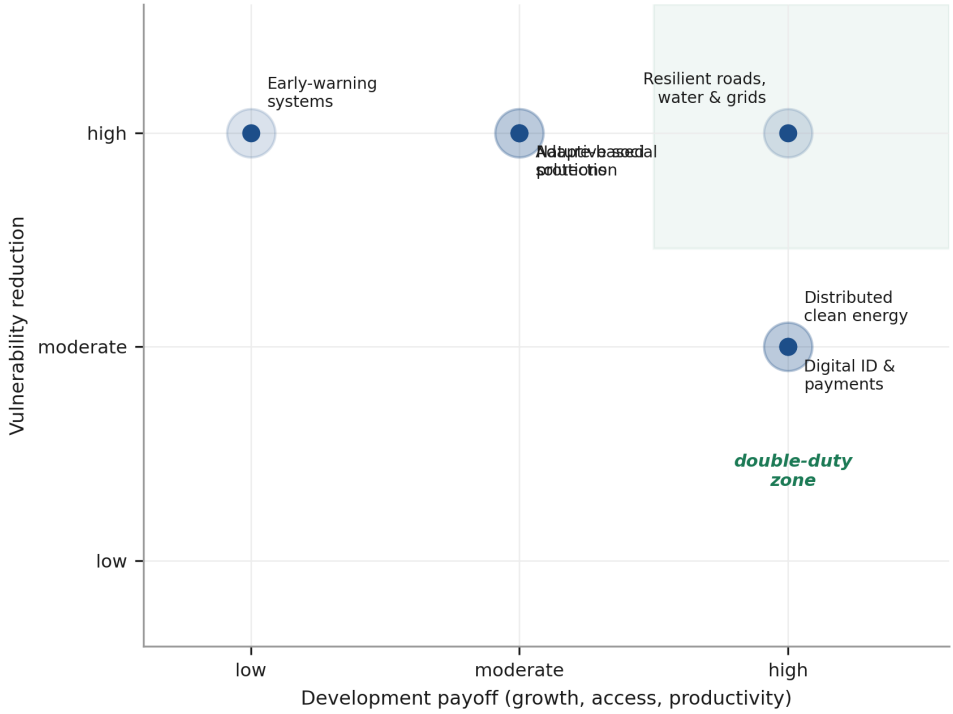

Figure 6. *Double-duty investments, scored by development payoff and vulnerability reduction.*
Illustrative positioning by the author, drawing on the sources cited throughout.

Table 2. Double-duty investments: strengthening development and resilience at once.

Investment	Development payoff	Vulnerability reduction	Illustrative example
Distributed clean energy	Cheaper, faster energy access; local jobs	Severs fuel-price channel; no single point of failure	Mini-grids; rooftop solar
Digital ID & payment rails	Financial inclusion; efficient service delivery	Enables cash to reach households within days of a shock	India's digital public infrastructure
Adaptive social protection	Poverty reduction; human-capital gains	Prevents distress asset sales after shocks	Kenya's HSNP
Early-warning systems	Protects economic activity and continuity	Averts mortality and asset loss; ~10:1 return	Bangladesh cyclone preparedness
Resilient transport, water & grids	Connectivity, health, productivity	Keeps services running through shocks	Climate-proofed roads and drainage

Investment	Development payoff	Vulnerability reduction	Illustrative example
Nature-based solutions	Fisheries, tourism, water security	Absorbs surge and flood; stores carbon	Mangrove restoration

The point of the double-duty lens is that each item can be funded on development grounds while also reducing vulnerability — lowering the political and fiscal bar for adoption.

The double-duty lens resolves a tension that otherwise stalls resilience policy. Ministries of finance are, rightly, reluctant to divert scarce capital from growth to a defensive purpose whose payoff is an avoided loss they may never see. Double-duty investments dissolve the trade-off: because they raise growth *and* cut vulnerability, they need not be sold as insurance. The resilience is a co-benefit of an investment that pays for itself on conventional grounds — which is why the resilient pathway, contrary to the intuition that it must be the slower one, can be the faster one as well.

7.2 The cost of pursuing them separately

The corollary is a warning. When adaptation and clean infrastructure are pursued in separate silos — different ministries, different funds, different plans — the synergies are lost and, worse, the levers can work against each other. A cooling programme rolled out on a fossil-heavy grid raises emissions and locks in demand; a hard flood defence that ignores upstream nature-based options can be more expensive and more brittle than an integrated design; a clean-energy build that ignores climate exposure can put new solar farms in the path of the next flood. Integration is not a bureaucratic nicety. It is the difference between the multiplicative returns of Figure 4 and a set of stranded, partial efforts that cost more and protect less.

8. The Financing Constraint and How to Loosen It

> *Everything above runs into the same wall: the resilient pathway is known, proven, and often cheaper in the long run, yet it is not being financed at anything like the required scale. The binding constraint is not knowledge or technology. It is capital — its cost, its availability, and the debt that pre-empts it.*

Three facts define the constraint. First, the needs are large: adaptation alone requires an estimated US$310–365 billion a year for developing countries by 2035 (Figure 7), and total climate-and-development investment needs run to several trillion dollars a year. Second, the cost of capital in developing economies is punishingly high — the 400-to-800 basis-point wedge of Figure 8 — which, because resilient and clean investments are capital-intensive and front-loaded, does more to block them than any other single factor. Third, existing debt pre-empts the fiscal space that might otherwise fund resilience: with US$921 billion a year going to interest and 3.4 billion people in countries that spend more on debt service than on health or education, the money is committed before the conversation about resilience even begins.

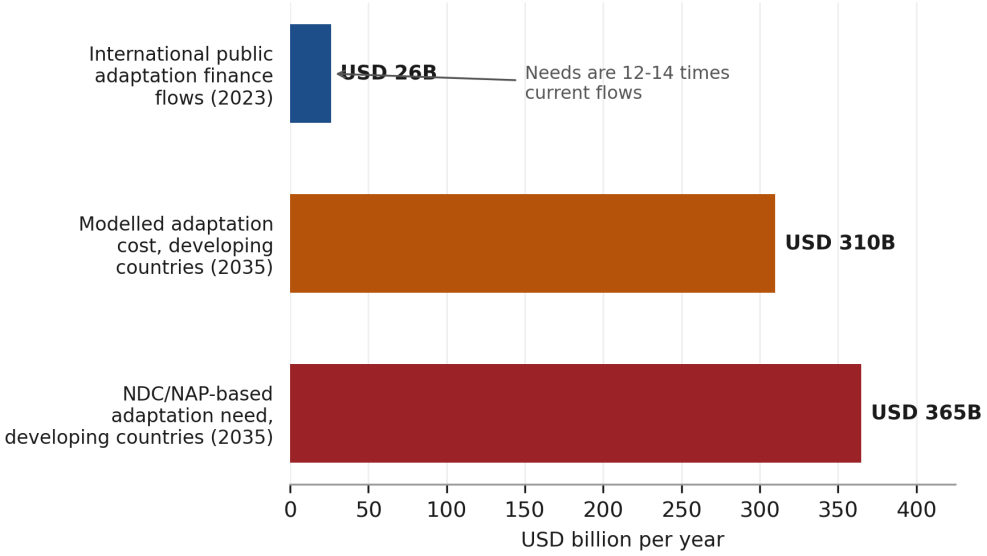

Figure 7. *The adaptation finance gap: needs versus international public flows.*

UNEP Adaptation Gap Report 2025 (needs of US$310–365 billion/year by 2035; international public adaptation flows of ~US$26 billion).

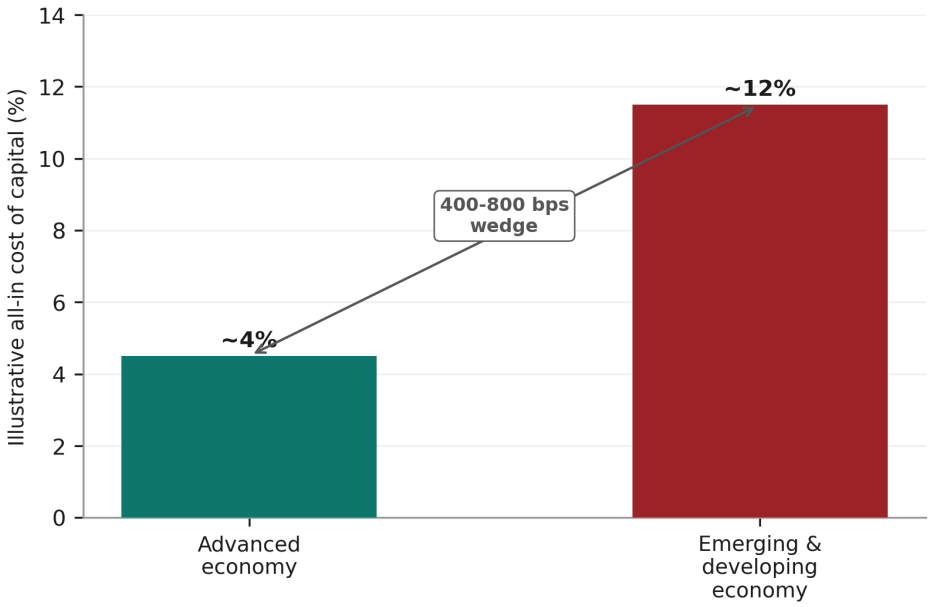

Figure 8. *The cost-of-capital wedge between advanced and developing economies.*

Illustrative, after IHLEG and LSE Grantham Institute analyses; the 400–800 bps wedge is widely documented and context-specific.

8.1 The cost of capital is a policy variable

The most important insight for policy is that the cost-of-capital wedge is not a fixed property of poor countries; it is substantially a function of how capital is structured, and therefore something policy can move. A large part of the premium reflects real and perceived risks — currency risk, off-taker and regulatory risk, political risk, and unfamiliarity — that can be reduced or reallocated by instruments designed for the purpose: partial guarantees, first-loss tranches, currency-hedging facilities, and long-tenor concessional finance that crowds in private capital rather than substituting for it. When these instruments lower the perceived risk, they lower the price of money, and because clean and resilient assets are so sensitive to that price, a modest reduction in the wedge unlocks a disproportionate volume of investment. De-risking is, in effect, the highest-leverage climate-and-development finance policy available.

This is the proper role of public and concessional finance, and of the multilateral development banks in particular. Their comparative advantage is not to fund the entire need from their own balance sheets — they cannot — but to absorb the specific risks that keep private capital out, and so to mobilise multiples of their own capital. The ongoing agenda of MDB reform — better use of existing balance sheets, more callable capital, more risk transfer, and a mandate that explicitly values mobilisation — is aimed at exactly this, and it is among the highest-return institutional reforms on the international agenda because it works on the price of capital rather than only its quantity.

8.2 Debt is the channel that must be cleared

No amount of new resilience finance will help a country whose fiscal space is entirely consumed by servicing old debt. Clearing the debt channel is therefore a precondition for the rest of the agenda, not a separate item. Three approaches matter. **Debt relief and restructuring** — faster, deeper, and more predictable than the current process — frees fiscal space directly. **State-contingent debt instruments**, such

as the natural-disaster and hurricane clauses pioneered by Barbados, pause repayments automatically after a catastrophe, converting debt from a shock amplifier into a shock absorber. And **liquidity support** through instruments like the IMF's Resilience and Sustainability Trust provides the affordable, long-tenor finance that lets a country invest in resilience without a fresh solvency crisis. These are not acts of charity; they are the clearing of a channel that otherwise transmits every shock into an unrecoverable one.

8.3 The emerging architecture

The international system has begun, haltingly, to respond. The 2025 Sevilla Commitment — the outcome of the fourth Financing for Development conference and the first such framework since 2015 — set out an agenda spanning debt, MDB reform, and domestic resource mobilisation. The COP process has raised, on paper, both the headline finance goal and the specific commitment to adaptation, with a call to triple adaptation finance toward roughly US$120 billion a year by 2035 and a roadmap toward US$1.3 trillion a year in total climate finance. The gap between these commitments and delivered flows remains vast, and this report takes no view on whether the pledges will be honoured. The point is narrower and more durable: the instruments that would loosen the financing constraint now exist and are understood. What is missing is scale — and, as the next section argues, the institutional capacity to absorb finance and turn it into resilience on the ground.

9. Institutions: Turning Money and Technology into Resilience

> *Capital and technology are necessary but not sufficient. The record is full of well-funded projects that failed to reduce vulnerability and of poor countries that weathered shocks far better than their wealth predicted. The difference is almost always institutional capacity — the least glamorous and most decisive lever of all.*

If finance is the constraint on the input side, institutions are the constraint on the conversion side: the machinery that turns money and technology into resilience actually delivered. This machinery is unevenly distributed and hard to build quickly, which is why it deserves explicit attention rather than being assumed. Three functions do most of the work.

- **Anticipatory capacity** — the ability to see a shock coming and act before it lands. This is the early-warning system of Section 5 extended into government: risk monitoring, contingency plans activated by pre-agreed triggers, and pre-arranged finance released automatically. Anticipatory action is far cheaper and more effective than post-disaster response, but it requires the standing institutional plumbing — the data, the triggers, the authority to act — to be in place in advance.
- **Delivery capacity** — the ability to plan, procure, build, and maintain. Much of the resilience gap is a project-preparation and maintenance gap: countries lack the pipeline of bankable, well-designed projects that finance can flow into, and lack the maintenance systems that keep resilient assets resilient over their long lives. A resilient bridge that is not maintained is, within a decade, just a bridge.
- **Data and coordination capacity** — the ability to know where the exposure and vulnerability actually are, and to coordinate across the many ministries whose decisions jointly determine the pathway. Resilience is cross-cutting by nature; without a coordinating function above the individual sectors, the silo problem of Section 7 reasserts itself and the synergies are lost.

9.1 The enabling infrastructure of resilience

Some institutional investments are enabling in the strong sense that they raise the return on everything else. Digital public infrastructure — a foundational identity system, an interoperable payments rail, and the data registries built on them — is the clearest case. It is what allows adaptive social protection to reach the right households within days, allows subsidies and relief to be targeted rather than sprayed, and allows an early warning to translate into an actual transfer of cash or an actual evacuation. A country that has built this plumbing can deploy a new resilience programme on top of it in months; a country that has not must build the plumbing first every time. This is why digital public infrastructure appears among the double-duty investments of Section 7: it is simultaneously a development accelerator and a resilience multiplier.

The broader lesson from the country evidence in Section 11 is that the states which absorb shocks well are distinguished less by their wealth or even their exposure than by whether they built and protected this enabling capacity — fiscal space, delivery institutions, pre-arranged finance, durable coordination — *before* the shocks arrived. Capacity built in advance is worth a large multiple of the same capacity assembled in the middle of an emergency. The implication for sequencing is direct, and it is the subject of the next section: the institutional and financial enablers should come first, because they determine how much of everything else can actually be delivered.

10. Sequencing: A Practical Order of Operations

> *A strategy that lists everything worth doing is not, in a world of binding constraints, a strategy at all. The value of the framework in this report is that it also implies an order — a defensible answer to the question of what a constrained government should do first.*

The sequencing principle follows from two facts established above. First, some investments are no-regret: they pay for themselves on development grounds regardless of how the climate and shock environment evolve, so they should be made immediately and unconditionally. Second, some investments are enablers whose main value is to raise the return on later ones — the finance and institutions of Sections 8 and 9 — so they should come early even when their direct payoff is modest, because they determine how much of the rest can be delivered. Everything else can then be phased against capacity and finance as they grow. Table 3 sets out the resulting three tiers.

Table 3. A three-tier sequencing rule for a constrained government.

Tier	Principle	Representative measures	When to move to the next tier
Tier 1 — No-regret & enabling	Do now, unconditionally	Early warning; adaptive social protection; digital ID & payments; project-preparation and data capacity; clearing the debt channel	Immediately; these are the foundation
Tier 2 — Enabling-contingent	Do as finance and capacity allow	De-risked clean-energy build-out; resilient core infrastructure; nature-based solutions; pre-arranged disaster finance	As Tier 1 institutions can absorb and maintain them
Tier 3 — Transformational	Phase in against locked-in enablers	System-wide energy transition; managed retreat and relocation; major structural diversification	Once financing and delivery capacity are durable

The tiers are about timing and conditionality, not importance; a transformational investment may be the most important thing a country does, but it should be sequenced after the enablers that make it deliverable.

Three features of this rule are worth drawing out. It puts the cheapest, highest-return, most robust measures first, so that even a government that can do only a little does the most valuable little. It front-loads the enablers — finance and institutions — because they are the multiplier on everything downstream, which means the early years of a resilient pathway are dominated by unglamorous plumbing rather than headline megaprojects. And it treats the transformational investments not as things to be deferred indefinitely but as things to be *prepared* for, so that when the enabling capacity is in place the country can move quickly rather than starting from scratch. The rule is conservative in sequence and ambitious in destination.

The rule also guards against a common failure mode: pouring finance into Tier 3 megaprojects before the Tier 1 enablers exist to deliver and maintain them. A large clean-energy or flood-defence programme launched without project-preparation capacity, maintenance systems, or coordinating institutions tends to underperform, stall, or decay — converting scarce capital into stranded ambition. Sequencing is not a counsel of caution; it is a counsel of order, and the order is what makes the ambition achievable.

11. Country Pathways: Evidence from Practice

> *The argument of this report is not theoretical. Countries at every income level have already bent their vulnerability curves using one or both levers, and their experience is the best available evidence for what the framework claims — and for its limits.*

The cases below are chosen not because any is a finished model but because each isolates part of the argument: the mortality dividend from anticipatory institutions, the safety net that scales on cue, the clean build-out that severs the fuel channel, the debt instrument that absorbs shocks, and the coordinating institution built in advance. Read together, they make the report's central empirical point: the countries that absorb shocks well are distinguished less by wealth or exposure than by whether they built the enabling capacity before the shocks arrived.

11.1 Bangladesh: the mortality dividend of anticipatory institutions

Bangladesh is the clearest demonstration that adaptive capacity, not the hazard, determines the human cost of disaster. The 1970 Bhola cyclone killed on the order of several hundred thousand people; a comparably powerful cyclone in 1991 killed around 138,000. Yet in recent years cyclones of similar intensity have killed in the tens or low hundreds. The hazard did not diminish — if anything it intensified — but Bangladesh built, over decades, a Cyclone Preparedness Programme combining a dense network of shelters, a volunteer early-warning system reaching the last mile, and evacuation protocols rehearsed in advance. This is the mortality line of Figure 2 made concrete, and it was achieved by a low-income country through institutions rather than wealth. Bangladesh has since extended the same anticipatory logic to floods, piloting forecast-based cash transfers that reach households *before* the water arrives — adaptation acting on adaptive capacity, exactly as the framework predicts.

11.2 Kenya: a safety net that scales with the shock

Kenya's Hunger Safety Net Programme shows adaptive social protection working as designed. Built on a registry of vulnerable households in the drought-prone north and a mobile-payment rail, the programme can expand automatically when a satellite-based drought index crosses a pre-agreed threshold, delivering cash to affected families within days. The speed is the point: it is the interval between the shock and the response that decides whether a household copes by drawing on the transfer or by selling livestock and pulling children from school. Kenya's system converts a slow, discretionary, post-hoc relief process into a fast, rules-based, anticipatory one — and in doing so illustrates why digital public infrastructure (the registry and the payment rail) is the enabling investment that makes the resilience possible.

11.3 Morocco and Vietnam: severing the fuel channel

On the clean-infrastructure lever, Morocco and Vietnam illustrate two routes to the same resilience gain. Morocco, an energy importer historically exposed to fuel-price shocks, invested heavily in large-scale solar and wind — including the Noor Ouarzazate complex — deliberately reframing renewables as an energy-security and balance-of-payments strategy as much as a climate one. Vietnam demonstrates the speed at which the clean pathway can be taken when the incentives align: a policy shift unleashed one of the fastest solar build-outs in the world in the space of a few years, adding many gigawatts of capacity, much of it

distributed rooftop generation. Both cases show clean infrastructure cutting the imported-fuel transmission channel of Section 3 — and both also show the coordination challenge, as rapid build-outs strained grids that had not been upgraded in step, a reminder that the levers must be integrated rather than pursued in isolation.

11.4 Barbados: making debt a shock absorber

Barbados has pioneered the financial side of resilience. Facing the classic small-island bind — high exposure to hurricanes, a narrow economy, and a heavy debt load — it pioneered natural-disaster and "hurricane" clauses in its debt instruments that automatically pause repayments after a qualifying catastrophe, freeing cash for response exactly when it is most needed and converting debt from a shock amplifier into a shock absorber. Through the associated Bridgetown Initiative, Barbados has pressed the same logic onto the international agenda, and it has drawn on the IMF's Resilience and Sustainability Trust for affordable long-tenor finance. The case shows that some of the highest-return resilience investments are contractual rather than physical — a clause in a bond can be worth a sea wall — and that a small state can shape the global architecture as well as use it.

11.5 Rwanda: coordination built in advance

Rwanda illustrates the institutional lever. Rather than leaving climate and resilience spending scattered across ministries, it established a national environment and climate fund (FONERWA) to pool, coordinate, and channel finance toward priority investments, backed by a broader reputation for delivery capacity and low leakage. The specific projects matter less than the institutional form: a standing, coordinating vehicle that can absorb finance, prepare projects, and direct them coherently is precisely the delivery-and-coordination capacity of Section 9. It is the kind of enabler that, once built, raises the return on every subsequent dollar of adaptation and clean-infrastructure finance — the capacity built in advance that the framework identifies as the true differentiator.

11.6 What the cases share

Across five very different countries the same pattern recurs. In each, the decisive investment was made *before* the shock it addressed, and in each it worked on a specific determinant of vulnerability rather than on the hazard: adaptive capacity in Bangladesh and Kenya, exposure and the fuel channel in Morocco and Vietnam, the financial transmission channel in Barbados, and coordination capacity in Rwanda. None is wealthy; several are among the poorer countries in their regions. What distinguishes them is not resources but foresight and institutional design — which is the most encouraging finding in this report, because foresight and design are available to any government willing to sequence its choices deliberately.

12. Measuring Progress and Avoiding Maladaptation

> *A resilience strategy that cannot tell whether it is working will not survive contact with a budget process, and a resilience investment that is poorly designed can raise vulnerability rather than lower it. Both problems — measurement and maladaptation — deserve explicit attention.*

12.1 Measuring an avoided loss

The central measurement difficulty in this field is that the payoff of resilience is a loss that does not happen. When an early-warning system works, the deaths it prevents are never counted; when a resilient road survives a flood, the disruption it avoids is invisible. This makes resilience chronically hard to justify against investments whose returns are observable, and it biases budgets toward the visible and away from the preventive. The response is not to abandon measurement but to shift it from outputs to the determinants of vulnerability: to track exposure (how many people and how much value sit in high-hazard areas), sensitivity (how much a standard shock damages key systems), and adaptive capacity (coverage of early warning and social protection, availability of pre-arranged finance, project-preparation and maintenance capacity). These are measurable in advance, they move in response to the two levers, and they are leading indicators of avoided loss rather than lagging counts of realised damage.

Good measurement also disciplines the strategy against its own worst tendency — the temptation to fund the visible and the new over the invisible and the maintained. Maintenance of existing resilient assets, replenishment of contingency finance, and the upkeep of registries and warning systems are exactly the unglamorous, preventive spending that budgets cut first and regret later. A measurement system anchored on the determinants of vulnerability keeps these on the ledger by making their erosion visible before the next shock does.

12.2 Maladaptation: when resilience spending backfires

Not all adaptation reduces vulnerability, and some increases it. The literature on maladaptation catalogues the failure modes, and a serious strategy must design against them. A hard flood defence can create a false sense of security that draws more people and assets into the protected floodplain, raising the eventual loss when the defence is finally overtopped. Irrigation that stabilises yields in the short run can deplete an aquifer and deepen vulnerability in the long run. An adaptation that protects one community by displacing water or risk onto its neighbour reduces aggregate vulnerability not at all. And a resilience measure built on a fossil-heavy base — air-conditioning on a coal grid, desalination powered by imported diesel — can lock in emissions and fuel exposure even as it addresses a local hazard, working against the very hazard reduction the clean lever is meant to deliver.

The common thread in these failures is a narrow frame: an intervention optimised for one hazard, one community, or one time horizon, without regard to the couplings and the long lives that Sections 2 and 4 insist upon. The guard against maladaptation is therefore the same integration this report has argued for throughout — assessing each investment against the whole anatomy of vulnerability and over the full life of the asset, and pursuing the two levers together so that an adaptation measure is not quietly undermining the hazard reduction, or vice versa. Maladaptation is, in the end, what the siloed pursuit of resilience looks like in practice, and it is the strongest practical argument for the integrated pathway.

13. Conclusion: Choosing the Pathway

> *The polycrisis is not a forecast of doom and it is not a reason for paralysis. It is a description of a risk environment in which shocks arrive together and travel between systems — and it carries within it a clear implication for how development should now be done.*

This report has made four connected claims. First, that the danger of the present period lies in the coupling of shocks, not their number, and that the coupling falls hardest on developing economies because that is where the transmission channels converge. Second, that the vulnerability which turns coupled shocks into systemic crises is built rather than given, and can therefore be un-built — by acting on its determinants of exposure, sensitivity, and adaptive capacity. Third, that two levers do most of that work: adaptation, which raises adaptive capacity and lowers sensitivity, and clean infrastructure, which cuts exposure, severs the fuel channel, and shrinks the hazard itself. And fourth, that these levers are complements whose combined return exceeds the sum of their parts, concentrated in a set of double-duty investments that pay in development terms and in resilience at once.

The constraints are real but they are not the ones usually assumed. The binding limits are not technological — the technology is proven and, for clean infrastructure, now often the cheapest option — nor are they a lack of knowledge about what works, which is unusually well established for adaptation. The binding limits are the cost and availability of capital, the debt that pre-empts fiscal space, and the institutional capacity to convert finance into resilience on the ground. Each of these is a policy variable, not a fact of nature: the cost of capital responds to how risk is structured, the debt channel can be cleared, and institutional capacity can be built — provided it is built in advance.

That last phrase is the whole argument compressed. Because infrastructure is long-lived, because vulnerability accumulates, and because capacity assembled in an emergency is worth a fraction of the same capacity built beforehand, the timing of the choice is everything. The window in which a developing economy decides the composition of its mid-century capital stock is open now and will not stay open; the premium on the resilient option is smallest before the concrete is poured. The countries examined here that weathered their shocks did not do so because they were rich or unexposed. They did so because they chose, earlier, to build the enabling capacity — and were therefore standing on a different pathway when the shock arrived.

The task for the coming decade is correspondingly concrete. It is not to invent a resilience agenda; the agenda exists, its returns are documented, and its instruments are understood. It is to fund that agenda by loosening the financing constraint, to integrate its two levers rather than pursuing them in silos, to sequence it so the enablers come first, and above all to build it in advance of the shocks it is meant to withstand. A developing economy cannot choose its exposure to a warming, shock-prone world. It can choose the pathway it takes through it — and that choice, made now, is the difference between development that is repeatedly reversed and development that endures.

References and Further Reading

> *Sources are grouped by theme. Every quantitative claim in the report is attributed inline to one of the works below; URLs point to the publishing institution's landing page for each work.*

Systemic risk, hazards, and the disaster record

- Intergovernmental Panel on Climate Change (IPCC), *Sixth Assessment Report, Working Group II: Impacts, Adaptation and Vulnerability* (2022). ipcc.ch/report/ar6/wg2
- World Meteorological Organization, *Atlas of Mortality and Economic Losses from Weather, Climate and Water Extremes (1970–2021)*. wmo.int
- UN Office for Disaster Risk Reduction (UNDRR), *Global Assessment Report on Disaster Risk Reduction*. undrr.org/gar
- World Economic Forum, *Global Risks Report 2026*. weforum.org

Adaptation: returns, gaps, and early warning

- Global Commission on Adaptation, *Adapt Now: A Global Call for Leadership on Climate Resilience* (2019). gca.org
- World Bank, *Lifelines: The Resilient Infrastructure Opportunity* (2019). worldbank.org
- UN Environment Programme, *Adaptation Gap Report 2025*. unep.org
- World Meteorological Organization / UN, *Early Warnings for All*. wmo.int/activities/early-warnings-all
- World Bank, *State of Social Protection Report 2025*. worldbank.org

Clean infrastructure and the energy transition

- International Energy Agency, *World Energy Investment 2025*. iea.org
- International Energy Agency, *SDG7: Data and Projections*. iea.org
- World Bank & African Development Bank, *Mission 300 (Energizing Africa)*. worldbank.org/en/programs/energizing-africa

Debt, cost of capital, and development finance

- UNCTAD, *A World of Debt 2025*. unctad.org/publication/world-of-debt
- Independent High-Level Expert Group on Climate Finance (IHLEG), reports on financing climate action and the cost of capital, LSE Grantham Research Institute. lse.ac.uk/granthaminstitute
- International Monetary Fund, *Resilience and Sustainability Trust*. imf.org
- United Nations, *Sevilla Commitment — Fourth International Conference on Financing for Development (FfD4, 2025)*. financing.desa.un.org
- Bridgetown Initiative. bridgetown-initiative.org

A note on figures and estimates

Figures 1 (the anatomy of vulnerability), 5 (the two development pathways), and 6 (double-duty positioning) are conceptual schematics that illustrate the report's framework rather than empirical series.

Figure 2 is an illustrative decadal series reflecting the documented trends in the WMO Atlas. Figures 4 (benefit–cost ranges) and 8 (the cost-of-capital wedge) present indicative ranges drawn from the cited sources, and Figure 3 uses standard engineering lifetime ranges. Figure 7 (the adaptation gap) reproduces values reported by the source named in its caption. All benefit–cost ratios are context-specific and are presented as ranges; the report is analytical rather than predictive, and no figure should be read as a forecast.