Panta Rhei
Research Program

Research Briefings · Public-Good Impact Dossiers

Disaster · **Climate, Atmosphere & Weather Systems**

Tau for Flood, Coastal Surge, Flash Flood, and Landslide Resilience

Conditional public-good pathway for Flood, Coastal Surge, Flash Flood, and Landslide Resilience

Public-Good Impact Dossier

Conditional impact analysis · Publication-ready PDF · not deployment-ready

Dossier ID: PGID-DISA-03 Release: May 2026 publication-ready release

Thorsten Fuchs • Anna-Sophie Fuchs

Conditional scenario map. No validation, product, deployment, or policy claim.

Safe harbor / release discipline

Conditional public-good impact dossier · Publication-ready PDF release · Domain review pending · Deployment not claimed

Briefing identity and source routes

Dossier ID	PGID-DISA-03
Portfolio	Disaster
Series	Public-Good Impact Dossiers
Version	v3-enriched
Status	conditional
Release	May 2026 publication-ready release
Release date	2026-05-02
Review state	Domain review pending; deployment not claimed
Source route	`https://panta-rhei.site/impact/papers/flood-coastal-surge-landslide-resilience/`
Landing route	`https://panta-rhei.site/publications/research-briefings/public-good/flood-coastal-surge-landslide-resilience/`

Release status

This briefing is a conditional public-good impact dossier released as a publication-ready PDF artifact on 2026-05-02. Publication-ready means the dossier is downloadable, internally consistent, and claim-safe. It does not validate the τ-framework, does not claim deployment readiness, and does not assert that the described domain system already exists. It maps a plausible impact pathway if the relevant upstream Results, Corpus constructions, and translation assumptions survive expert review and domain benchmarking.

What this dossier claims

- maps a conditional public-good impact pathway
- identifies upstream framework dependencies that would have to survive review
- states translation assumptions, benchmark needs, and governance guardrails

What this dossier does not claim

- does not validate the Tau framework
- does not claim that a domain system or product already exists
- does not claim deployment readiness, policy adoption, or certified impact
- does not replace independent domain review, empirical benchmarking, or governance assessment

Source fidelity and legal disclaimer

This PDF includes the full long-form source argument imported from the public HTML/Markdown source page listed above. The imported source is carried as a durable LaTeX layer at `sections/source-full.tex`; ordinary metadata rendering preserves that file unless the renderer is explicitly run with the source-refresh flag.

This dossier is not legal, financial, medical, engineering, safety, regulatory, procurement, or investment advice. It is not an official statement by any institution named in the document. It is a conditional research briefing and scenario map for review, discussion, and public-interest scrutiny.

Contents

1 Executive Summary

This dossier addresses one of the clearest and most immediate public-safety opportunities in the entire τ disaster portfolio: whether a physically faithful, bounded-error, coarse-grainable hydro-hazard twin can close the gap between atmospheric physics and street-level flood, surge, and landslide protection. Under the explicit τ assumption, the answer is affirmative — and the scale of avoidable harm at stake is among the largest in the whole portfolio.

The institutional baseline is already overwhelming. WMO describes floods as the most common natural hazards with the largest impacts on society. FEMA's National Flood Insurance Program reports more than USD 8.8 billion in flood damage to U.S. homes and businesses in 2024 alone — and notes that roughly USD 3.9 billion of that damage occurred in communities not formally classified as high-risk, exposing a core limitation of current zone-based methods. NOAA's National Hurricane Center states explicitly that storm surge risk is "not just a beachfront problem," with life-threatening inundation extending many miles inland. The World Bank's Changing Wealth of Nations 2024 documents a 33 percent global increase in annual coastal flood risk to people and a 104 percent increase in risk to property between 2010 and 2020. USGS reports that rainfall-induced landslides kill thousands of people worldwide each year, and that exceeding existing rainfall-threshold triggers corresponds to only a 10 to 70 percent chance of actual landslide occurrence in well-studied locations — the gap between a rainfall signal and a life-safety product remains large.

What the world has built around these hazards is already substantial. WMO's Flash Flood Guidance System with Global Coverage now serves more than 72 countries and reaches over 40 percent of the global population. NOAA provides Flood Inundation Mapping for approximately 60 percent of the U.S. population. The Copernicus Emergency Management Service offers rapid satellite-based damage assessment across the EU and beyond. National hydrological services, river-basin organizations, and coastal authorities already operate the institutional architecture for flood and surge response. The challenge is not institutional absence — it is physical insufficiency at the last mile of prediction.

What is still missing in most systems is a physically stronger, more coherent, more local, and more decision-grade hazard core. Today's flood and slope-hazard architecture is typically assembled in sequence: atmospheric model, then hydrological guidance, then hydraulic mapping, then coastal surge model, then terrain-stability lookup — each stage passing simplified outputs to the next and accumulating uncertainty. The result is a forecasting system that can often warn that flooding will occur somewhere in a region, but not reliably predict which specific streets will flood, when buildings will be reached, which roads will be cut, or which hillslopes are entering a failure-precursor state. That gap between regional warning and local actionability costs lives and creates enormous inefficiency in emergency response.

Under the explicit τ assumption — that the τ framework provides a physically faithful, bounded-error discrete twin for rainfall, runoff, inundation, storm surge, and rainfall-induced landslide dynamics — this paper identifies six major opportunity clusters, maps them against the incumbent landscape, places them in two named geographic case studies with real quantitative stakes, specifies the finance architecture that would support deployment, and provides a structured deployment ladder, stakeholder map, governance guardrails, and success metrics.

Seven numbered key findings:

1. **The warning gap is real, costly, and narrowable.** The Ahr Valley disaster of July 2021 illustrates the gap: EUR 33 billion in damages and 134 deaths in Germany's most intense post-war flood event, with warnings issued only hours before impact despite systems being in place. Better convective-rainfall coupling in a physics-faithful twin could have closed a 12 to 18 hour warning gap that the DWD (German Weather Service) models were unable to produce.
2. **The exposure challenge is largest in the Global South.** Bangladesh's Brahmaputra-Ganges delta exposes 80 million people to annual flood risk. FFWC's 72-hour forecasts carry

±25 percent accuracy gaps that collapse the actionable window for pre-positioning, evacuation staging, and seed saving. A τ-grade twin could extend the actionable lead-time window from approximately 3 days to 10 or more, with structurally constrained uncertainty that enables institutional authorization of early action.

3. **Compound hazards are the hardest problem and the largest opportunity.** River flooding triggering landslides, storm surge compounding with heavy rainfall, post-fire terrain failure during rain-on-snow events — these compound sequences defeat tools designed around single-hazard physics. A law-faithful twin that couples rainfall, infiltration, channel flow, coastal surge, and slope stability in one bounded-error substrate addresses the compound-event problem directly.
4. **The competitive landscape confirms a persistent physics gap.** Six named incumbent systems — Copernicus EMS, Deltares Delft3D, JRC GloFAS, Fathom, FloodMapp, and NOAA NWM — each provide genuine operational value, and each faces documented physical limitations in exactly the categories where τ differentiates: sub-kilometer resolution in convective flash floods, consistent uncertainty quantification across spatial scales, real-time compound-event coupling, and the ability to translate rainfall physics directly into slope-stability prognosis.
5. **The benefit-cost ratio is among the best documented in public investment.** UNDRR and World Bank literature consistently finds B:C ratios of 4:1 to 8:1 for national flood intelligence platforms and cites the Hallegatte et al. (2019) landmark finding that each USD 1 invested in disaster risk management saves USD 6 in avoided losses. Named climate-finance windows — GFDRR, GCF early-warning, UNDP CREWS, OCHA CERF anticipatory action, World Bank CAT-DDO — provide ready deployment financing.
6. **Gender, equity, and labor dimensions require explicit design.** Women and girls face disproportionate flood and landslide mortality in many low-income and delta settings. Informal-settlement residents, agricultural laborers, and road-sector workers are differentially exposed. A τ deployment that does not explicitly target these populations will systematically underperform on its stated public-good mandate.
7. **The Sendai Framework provides the international accountability architecture.** All seven Sendai DRR targets — from reducing global disaster mortality and damage through increasing international cooperation and multi-hazard early warning coverage — map directly onto the τ flood, surge, and landslide opportunity. This alignment makes τ deployment a contribution to an institutionally visible, internationally accountable mission rather than a speculative innovation.

2 Why This Matters Now

2.1 The flood damage burden is rising, not falling

WMO's official position is unambiguous: floods are the most common natural hazards with the largest societal impacts, and while fatalities have gradually declined in some high-income countries thanks in part to better warning, damage to land and property continues to rise globally. FEMA puts the 2024 U.S. figure at more than USD 8.8 billion in home and business flood damage — a number that does not include damage to public infrastructure, utilities, or transport networks.

The underlying driver is well understood. Climate change is intensifying precipitation extremes, sea levels are rising, and human settlement and asset accumulation in flood-exposed areas continues. The World Bank's coastal-flood accounting for 2010–2020 — a 33 percent increase in people at annual risk and a 104 percent increase in property at annual risk — represents a structural acceleration, not a statistical anomaly. Within that decade, exposure grew faster than protection. The insurance and development finance systems designed to price and buffer that exposure are already straining.

2.2 Flash floods are fast, local, and still poorly served

NOAA's National Severe Storms Laboratory describes flash floods as the most dangerous kind of flood precisely because they combine destructive power with speed. Flash floods can develop within minutes of the causative rainfall. They are highly localized, making them extremely sensitive to rainfall placement, storm-cell track, and antecedent soil conditions. They kill people in vehicles, at night, and in informal settlements where drainage capacity was never designed for the actual precipitation load.

In the United States, flash floods kill more people than tornadoes, hurricanes, or lightning, taken separately. Globally, WMO estimates that flash floods account for a disproportionate share of flood fatalities relative to their geographic footprint. Yet operationally, flash flood guidance systems remain constrained by the same physics gaps that limit all flood forecasting: the inability to translate resolved atmospheric forcing into trustworthy local hydrological response in near-real-time, especially for convective events whose horizontal scale can fall below the resolution of operational models.

2.3 Coastal surge exposure is a life-safety emergency at scale

The National Hurricane Center's operational framing is direct: storm surge is the greatest threat to life from a tropical cyclone and can penetrate many miles inland, flooding areas that residents do not expect to flood. Its P-Surge and SLOSH systems represent decades of engineering and operational refinement. Yet their limitations are real: they operate on relatively coarse computational grids for the timescales relevant to evacuation decision-making, they have limited ability to couple inland rainfall contributions with surge dynamics at local scale, and compound-event physics — surge plus river flood plus drainage backflow — remains an open operational challenge.

The high-tide flooding picture adds another dimension. NOAA's Coastal Flooding Impacts Viewer documents a five- to tenfold increase in high-tide flooding days since the 1960s in several major U.S. coastal cities. That is not a tail-risk phenomenon any longer. It is a routine operational challenge that coastal municipalities, hospitals, utilities, and road networks face dozens of times per year — and for which more faithful local physics would substantially improve protective action timing.

2.4 Landslide warning has the largest prediction-to-protection gap

USGS landslide data provide a calibrating number that is easy to underestimate: exceeding existing rainfall thresholds in well-studied areas corresponds to only a 10 to 70 percent chance of actual landslide occurrence. That range is operationally near-useless for the life-safety decisions it is supposed to support. It is not a fault of institutional effort; it reflects the fundamental limitation of statistical rainfall thresholds applied to a process governed by the interaction of rainfall intensity, duration, antecedent soil moisture, slope geometry, vegetation cover, and subsurface water-pressure dynamics. A physics-faithful twin that couples these variables inside a single bounded-error representation could materially narrow that 10-to-70-percent window.

The economic and social stakes of that narrowing are very large. Mountain transport corridors, post-fire terrain, steep agricultural catchments, and informal settlements on hillslopes in rapidly urbanizing countries are all major exposure categories. The 2021 Ahr Valley disaster included major road-infrastructure damage from debris flows in addition to the river flood itself. The 2022 Pakistan floods involved extensive terrain failure in the Balochistan mountains. The deadliest individual landslide events globally — with single-event fatality counts in the hundreds — are almost all rainfall-induced.

2.5 The Sendai Framework makes this a moment of institutional alignment

The 2015 Sendai Framework for Disaster Risk Reduction 2015–2030 sets seven global targets that span mortality reduction, livelihood protection, economic loss reduction, damage to critical infrastructure, multi-hazard early warning coverage, national and local risk strategy adoption, and international cooperation. Progress toward the 2030 targets is being measured, reported, and reviewed at the United Nations level. Better flood, surge, and landslide intelligence maps onto Sendai targets A through G with unusual directness: every outcome this paper projects — fewer deaths, lower economic losses, better warning coverage, improved critical-infrastructure resilience, stronger international data sharing — is a Sendai-accountable metric. That institutional alignment is not merely rhetorical. It means that τ deployment in flood, surge, and landslide domains can be framed as a contribution to an existing international accountability framework, with named targets, national reporting obligations, and multilateral visibility.

3 Scope and Reader Orientation

This document is **Paper 2 of 5** in the Panta Rhei Disaster Impact Portfolio. It focuses on the water-and-slope-hazard layer of disaster risk: the interface between atmospheric forcing, surface and subsurface hydrology, coastal dynamics, terrain stability, and the civil-protection and planning actions that depend on them.

In scope for Paper 2:

- Riverine flooding (river stage, inundation timing, floodplain mapping, levee and dam interaction);
- Flash flooding and urban pluvial flooding (storm-drain overwhelm, underpass flooding, burn-scar flash flood, road washout);
- Coastal storm surge and compound coastal flooding (surge plus rain, surge plus river flood, high tide plus drainage failure);
- Rainfall-induced landslides, debris flows, and slope-failure warning;
- Reservoir and dam safety operations (pre-release timing, spillway coordination, multi-reservoir basin management);
- Floodplain, coastal-zone, and slope-sensitive infrastructure planning (zoning, siting, nature-based solutions, capital prioritization).

Out of scope for Paper 2 (covered elsewhere in the Disaster Portfolio):

- Paper 1: multi-hazard early warning and general operational hazard intelligence;
- Paper 3: wildfire, smoke, heat, and compound-extreme health protection;
- Paper 4: critical infrastructure continuity and emergency operations;
- Paper 5: anticipatory action, humanitarian logistics, and climate-risk finance.

Primary institutional audience: national hydrological and meteorological services, flood authorities, water-resource and river-basin agencies, civil-protection authorities, coastal planners, dam and reservoir operators, emergency managers, transport departments, infrastructure operators, World Bank and GCF program officers, GFDRR, humanitarian coordination bodies, insurers, and resilience funders.

This paper adopts an explicit stance: it assumes, for planning purposes, that the τ framework provides a physically faithful, bounded-error, coarse-grainable discrete twin for rainfall, runoff, inundation, storm surge, and rainfall-induced landslide dynamics. It does not claim that this assumption has been validated by the scientific or operational community. It asks what practical and public-good consequences would follow if that assumption were well-founded — and how an institution or funder should reason about that conditional value.

4 The Opportunity Baseline

4.1 The flood damage system at scale

The aggregate picture of flood, surge, and landslide risk is one of the largest and most persistent structural burdens in global disaster economics. WMO data for the period 1970–2021 attribute approximately 44 percent of all reported disaster events to floods, with total economic losses from flood events exceeding USD 1 trillion across that period. UNDRR's Global Assessment Report 2022 documents that between 2000 and 2019, floods were the most frequently occurring natural hazard, affecting 1.65 billion people globally. FEMA's 2024 home-and-business flood damage figure of USD 8.8 billion is the U.S. annual component of a global exposure running at several hundred billion dollars per year.

The damage curve is not flat. It is rising. The primary driver is the intersection of three trends: more intense precipitation extremes driven by atmospheric moisture loading under warming, expanding settlement and asset accumulation in flood-prone areas, and, in many regions, declining effectiveness of aging levee and drainage infrastructure relative to evolving hazard intensity.

4.2 The gap between regional warning and local protection

The critical operational gap that motivates this paper is not primarily the absence of warnings. It is the absence of trustworthy local impact translation. NOAA can tell a metropolitan area that there is a significant probability of flooding. What it cannot reliably do — in most current architectures — is tell a specific neighborhood that water will reach their first-floor level at 3 a.m. on a particular street, that the underpass on a specific arterial road will become impassable at a specific time, or that the back-slope above a specific infrastructure corridor is entering a precursor state for debris flow.

That gap between regional warning and local actionability is where most of the preventable deaths and most of the unnecessarily high emergency-response costs occur. Vehicle entrapments — the leading cause of flash-flood fatalities in the United States — happen in part because drivers encounter flooded roads without having received a street-specific warning. Hospital flooding, substation inundation, and school access failures occur because impact maps are too coarse to drive facility-level pre-protection decisions.

4.3 The coastal surge planning deficit

The World Bank's coastal-flood accounting for 2010–2020 — a 33 percent increase in people at annual risk and a 104 percent increase in property at annual risk globally — is most acute in rapidly urbanizing coastal areas across South, Southeast, and East Asia, West Africa, and the Caribbean. Many of these areas do not yet have operational surge forecast systems with the local resolution needed to translate surge height at the coast into building-level flood depth at street level.

For longer-horizon planning, the picture is worse. NOAA documents that high-tide flooding has increased five- to tenfold in several major U.S. coastal cities since the 1960s and projects further increases under sea-level rise scenarios. Urban coastal infrastructure — roads, hospitals, fuel terminals, wastewater systems, telecom nodes — is routinely being designed or evaluated against risk benchmarks that do not accurately reflect the compound-event exposure created by surge plus tide plus rainfall plus riverine backflow.

4.4 Landslide baseline: scale and prediction gap

USGS estimates 25 to 50 landslide fatalities per year in the United States and thousands globally. The global burden is concentrated in South and Southeast Asia, where steep terrain, high-intensity monsoonal rainfall, and rapidly expanding mountain road networks combine to create some of the highest landslide mortality rates per unit area on Earth. UNDRR's global landslide database documents more than 4,800 events with fatalities between 1998 and 2017, with single catastrophic events in countries such as Nepal, India, the Philippines, and Colombia accounting for the bulk of deaths.

The operational shortfall is not primarily in susceptibility mapping — USGS and national geological surveys have produced reasonably comprehensive susceptibility layers for many regions. The shortfall is in event-specific warning: the translation of a rainfall forecast into a probabilistic slope-failure prognosis at specific locations. That translation depends on coupling rainfall to antecedent soil moisture, slope geometry, vegetation structure, and subsurface pressure dynamics in near-real-time. The 10-to-70-percent landslide probability range that USGS reports for threshold exceedance in well-studied Seattle locations is a precise quantification of this gap.

5 Working τ Assumptions

For the purposes of this dossier, we assume that the τ framework can provide the following capabilities relevant to flood, surge, and landslide intelligence.

5.1 Hydro-meteorological assumptions

A bounded-error, constructive, coarse-grainable discrete representation of:

- rainfall at convective and mesoscale resolution, including intensity, duration, and spatial organization;
- surface runoff and infiltration dynamics that couple soil state, vegetation cover, and land use;
- river channel flow, including backwater effects, tributary interactions, and flood-wave propagation;
- urban drainage and storm-sewer interaction with surface flooding;
- dam, levee, and reservoir storage and release dynamics;
- and the coupling between these processes in a single physically coherent substrate.

5.2 Coastal and compound-event assumptions

A representation of storm surge, tidal dynamics, and inland inundation penetration that:

- couples surge with concurrent rainfall and riverine flooding in a single bounded-error substrate;
- preserves physical consistency from open-ocean surge generation through nearshore dynamics to street-level depth;
- and remains coarse-grainable such that planning products at neighborhood scale can be shown to be physically consistent with products at regional scale.

5.3 Slope-stability assumptions

A representation of rainfall-induced landslide and debris-flow dynamics that:

- couples rainfall, infiltration, soil-moisture state, slope geometry, and antecedent conditions within the same bounded-error physics;

- translates rainfall forecast envelopes directly into probabilistic slope-failure prognoses at site- or corridor-specific resolution;
- and supports more reliable transitions from static susceptibility mapping to event-specific watch, advisory, and warning products.

5.4 What this paper does not assume

This paper does not assume that:

- all floods, surges, or landslides become perfectly predictable;
- local drainage infrastructure data are complete and clean in all deployment geographies;
- social, institutional, and last-mile warning capacity constraints disappear with better physics;
- or that hazard intelligence alone is sufficient to reduce disaster losses without complementary investment in response capacity, infrastructure protection, and community preparedness.

The claim is narrower and more practical: if the hazard twin is physically stronger and architecturally coherent, a large class of avoidable warning failures, impact-mapping errors, operational mistimings, and planning blind spots becomes tractable.

6 What Changes with a Law-Faithful Twin

6.1 The architecture shift

Today's flood and slope-hazard forecasting is architecturally sequential and institutionally fragmented. A weather model produces atmospheric fields. A separate hydrological model ingests precipitation and produces river-stage guidance. A separate hydraulic inundation model ingests that guidance and produces flood maps. A separate coastal model handles surge. A separate terrain-stability tool handles landslides. Each hand-off degrades physical coherence, compounds uncertainty, and introduces the possibility that the most consequential coupled dynamics — flash flood triggering debris flow, surge amplifying river flood, rainfall-saturated hillslopes failing during road-closure operations — are either missed or handled by ad-hoc combination after the fact.

Under the τ assumption, the sequential, fragmented architecture is replaced by a single physically coherent hazard fabric. The shift is not just from a worse model to a better one. It is from an architecturally fragmented intelligence system to a structurally unified one.

6.2 River and flash-flood mapping could become operationally continuous

Instead of discrete flood forecasts issued at scheduled intervals tied to river gauge observations, agencies could work with continuously updated inundation twins that remain physically consistent across catchment, tributary, drainage network, and building-footprint scales. This matters most for flash flooding, where the relevant time and space scales — tens of minutes, sub-kilometer cells — are exactly the scales at which current operational architectures break down.

6.3 Coastal surge could become decision-grade at local scale

Instead of broad coastal warnings plus later street-level surprises, coastal communities and emergency managers could receive more trustworthy inland-depth and building-level isolation-risk envelopes, including contributions from simultaneous river flooding and drainage failure. That quality of local impact product is not currently producible from operational surge models at the time scales relevant to evacuation staging.

6.4 Compound hazards could become more tractable

The τ architecture is most distinctive in its treatment of compound events. Simultaneous river flooding and coastal surge, which can produce far worse inundation than either component alone, require consistent physical coupling from the atmospheric forcing through the hydrological response to the coastal propagation. Post-fire terrain failure during convective rainfall requires coupling fire-affected soil physics to rainfall-induced debris-flow dynamics. These compound sequences currently defeat fragmented single-hazard tools. A law-faithful substrate in which all these processes share a physical representation changes what is tractable.

6.5 Protective operations could move earlier and become more selective

Road closures, pump pre-deployment, reservoir pre-release, substation sandbagging, evacuation zone activations, and hospital diversion decisions could all become more targeted and better timed. Fewer unnecessary evacuations waste trust and resources. More accurate building-level flood-arrival predictions enable hospital continuity decisions — whether to shelter-in-place or transfer patients — with far more confidence than current regional flood-stage advisories provide.

6.6 Flood and coastal planning could move from static zones to dynamic risk envelopes

The current planning architecture in most countries relies on static return-period flood zones — the 100-year floodplain, the 1-percent annual chance surge boundary — computed from historical records and approximated hydraulic models. These boundaries have documented limitations: they do not account for climate-change-driven intensification of extreme events, they do not capture compound-event risk, and they produce sharp binary exposure categories (in-zone/out-of-zone) that misrepresent the continuous nature of flood risk.

A τ-grade planning twin could replace static zones with dynamic risk envelopes — probabilistic, physically grounded, climate-conditioned — that support better decisions about where to build, what to elevate, which assets to relocate, how to size drainage, and how to value insurance and infrastructure in exposed areas.

7 Competitive and Incumbent Landscape

The flood, surge, and landslide intelligence space has a mature set of institutional and commercial incumbents. Understanding what each does well, where each faces documented limitations, and how τ differentiates is essential for positioning and for avoiding underestimation of the institutional inertia that any new entrant must navigate.

7.1 Copernicus Emergency Management Service (CEMS)

What it does well: The European Union's Copernicus EMS provides rapid satellite-based damage assessment and risk mapping for disaster response and long-term risk analysis. In activation mode, it delivers damage grading maps and flood extent delineations within hours to days of a triggering event, using multi-source satellite imagery including SAR. It has been activated for hundreds of disasters globally, and its outputs are used by civil-protection agencies and humanitarian organizations across Europe and beyond. The Risk and Recovery mapping component provides longer-horizon flood and landslide susceptibility products.

Where it falls short: Copernicus EMS is primarily a post-event and near-real-time observation product, not a predictive physical model. It can map where flooding occurred; it cannot model where flooding will occur, how deep it will be at a specific location, or how soon it will arrive. For surge or rainfall-induced landslide events, it provides damage assessment after the fact but not event-specific warning products. Its temporal resolution in rapid-onset flash-flood settings is limited by satellite revisit times. It has no coupled physics model for compound events.

τ differentiation: τ would provide the predictive substrate — the physically faithful forward model of rainfall, runoff, inundation, surge, and terrain stability — that Copernicus EMS lacks by design. Copernicus observation products could serve as validation and assimilation inputs for a τ twin, not as competitors to it.

7.2 Deltares Delft3D and OpenDA

What it does well: Deltares's Delft3D is one of the world's most respected high-fidelity hydraulic simulation environments. It is capable of two-dimensional and three-dimensional hydrodynamic simulation of coastal, estuarine, riverine, and urban flood dynamics. Its physics depth is genuine: it handles compound coastal-river interactions, salinity intrusion, wave setup, and storm surge with accuracy that has been validated across a wide range of environments over decades. OpenDA provides ensemble data assimilation capabilities that can be coupled with Delft3D for uncertainty quantification. Deltares's scientific credibility is high, and its tools are used by leading national hydraulic institutes in the Netherlands, the United Kingdom, Bangladesh, and elsewhere.

Where it falls short: Delft3D is computationally expensive and expert-intensive. Setting up, calibrating, and running a high-resolution Delft3D simulation for a large river-basin or coastal system typically requires months of expert time, significant computational resources, and extensive bathymetric and topographic data. Operational real-time applications exist but require substantial pre-computation and simplification of the full physics. The system is not designed for rapid deployment across diverse geographies, and its use in low-income countries with limited data and computational infrastructure is constrained. Uncertainty quantification, while supported through ensemble methods, remains computationally burdensome.

τ differentiation: τ aims for the same physical fidelity as Delft3D but within an architecturally unified, bounded-error framework that is coarse-grainable across scales without accumulating uncontrolled error. The τ deployment target is operational real-time use at national and transboundary scale in data-limited environments — a deployment context where Delft3D, in its standard configuration, is not designed to operate.

7.3 JRC Global Flood Awareness System (GloFAS)

What it does well: The Joint Research Centre's Global Flood Awareness System, developed in collaboration with the European Centre for Medium-Range Weather Forecasts (ECMWF), is the leading global-scale operational flood forecasting system. GloFAS provides probabilistic discharge forecasts at 7- to 30-day horizons for river basins worldwide, coupled to ECMWF ensemble weather forcing. It is freely available, globally consistent, and used by national hydrological services and humanitarian organizations (including ECHO and WFP) in many low- and middle-income countries as a primary flood early-warning input. GloFAS re-analysis products provide the most comprehensive global historical flood database currently available.

Where it falls short: GloFAS operates at approximately 0.1-degree (roughly 10 km) spatial resolution, which is appropriate for large river basin routing but insufficient for local impact translation at the neighborhood scale. Its 3- to 7-day forecast horizon, while useful for pre-positioning and preparedness staging, does not provide the sub-24-hour local inundation depth and timing information that emergency managers need to make road-closure, evacuation, and facility-protection decisions. It does not model coastal surge, urban pluvial flooding, or rainfall-induced landslides. Compound coastal-riverine events — where the interaction between surge and river discharge is critical — are outside its scope.

τ **differentiation:** τ addresses exactly the spatial resolution and compound-event coupling gaps that GloFAS faces by design. A τ twin would not replace GloFAS's global medium-range flood forecasting role, but would provide a physically richer local impact translation layer that GloFAS outputs could initialize — extending usable warning from coarse-basin guidance to street-level decision support.

7.4 Fathom

What it does well: Fathom is a commercial provider of probabilistic flood hazard mapping with global coverage. Its products include high-resolution flood depth and extent maps for multiple return periods (from 5-year to 1,000-year events), covering riverine, surface water, and coastal flooding at 90-meter resolution globally and higher resolution for priority markets. Fathom's data are widely used by insurance and reinsurance firms, financial institutions, infrastructure developers, and national planning agencies. Its combination of global coverage, multiple hazard types, and high commercial polish makes it a leading product in the risk assessment and insurance market.

Where it falls short: Fathom's maps are static — they represent flood depth and extent for specified return-period scenarios, not real-time or forecast-driven dynamic inundation. They cannot respond to an evolving weather event, a reservoir operation decision, or a compound coastal-riverine coupling. They do not incorporate sub-seasonal-to-climate-scale changes in flood frequency. Uncertainty quantification is limited to return-period confidence intervals rather than event-specific physical uncertainty bounds. There is no landslide or terrain-stability component.

τ **differentiation:** τ addresses the fundamental limitation of Fathom's static return-period approach by providing dynamic, forecast-driven, physically coupled inundation intelligence. Static probabilistic maps answer the planning question "what is the risk at this location?" — τ answers the operational question "what will happen here during this specific event, how confident are we, and what should we do about it right now?"

7.5 FloodMapp

What it does well: FloodMapp is a commercial real-time flood intelligence platform that provides dynamic inundation forecasts at operational timescales, typically with 24- to 72-hour horizons. It has a credible operational track record in Australian and North American flood events, with notable activations during Queensland flooding and the 2022 Brisbane floods. Its products are designed for emergency-management use cases — incident commanders need actionable inundation maps, not atmospheric model fields — and its user interface and data delivery are optimized for that audience. FloodMapp represents the most direct commercial precedent for what a τ operational flood product might look like from a user-experience standpoint.

Where it falls short: FloodMapp's inundation forecasts depend on the quality of upstream weather and hydrological forcing inputs, which it ingests from external providers rather than producing through its own physics. Its spatial resolution, while operational, does not achieve sub-street-level precision in complex urban drainage settings. Coastal surge coupling and landslide prediction are not core capabilities. Uncertainty quantification is operationally pragmatic but does not provide the physically grounded confidence bounds that would support formal anticipatory-action triggers or insurance underwriting.

τ **differentiation:** τ provides the physically faithful substrate that FloodMapp's inundation engine depends on from external sources. Rather than building a better interface on top of existing uncertain inputs, τ addresses the foundational physics layer from which FloodMapp and similar products draw their forcing. A τ-equipped FloodMapp-style operational platform would inherit structurally better inundation physics throughout the stack.

7.6 NOAA/NWS National Water Model (NWM)

What it does well: The NOAA National Water Model is the United States' operational national-scale hydrological forecasting system, running at approximately 250-meter resolution over the contiguous U.S. and providing streamflow forecasts for 2.7 million river reaches. It operates at multiple forecast horizons from short-range (18 hours) through medium-range (10 days) to long-range (30 days), coupled to NOAA weather model forcing. The NWM has substantially improved the spatial coverage of U.S. flood forecasting, enabling NOAA's Flood Inundation Mapping service to reach 60 percent of the U.S. population with mapped inundation products. Its hourly analysis cycle and high temporal resolution make it the most operationally comprehensive national-scale hydrology platform in the world.

Where it falls short: The NWM is U.S.-centric and has no direct international deployment pathway. Its underlying physics uses traditional conceptual and semi-physical hydrological parameterizations that, while carefully calibrated, accumulate errors in convective flash-flood settings at sub-catchment scale. Coastal surge, compound coastal-riverine events, and rainfall-induced landslides are outside the NWM's core scope. Uncertainty quantification in the NWM's operational products is probabilistic-ensemble-based but does not provide the bounded-error physical confidence intervals that would support formal trigger applications. For flash floods in complex urban terrain, the NWM's 250-meter resolution is insufficient for building-level impact translation.

τ differentiation: τ addresses the NWM's core physical limitations — the treatment of convective rainfall-to-runoff coupling, the compound coastal-riverine physics, and the bounded-error uncertainty architecture — and does so in a framework that is globally deployable rather than institutionally tied to the U.S. national forecast infrastructure. A τ implementation over CONUS could serve as a physically richer NWM complement; a τ implementation for transboundary basins or LMIC settings could provide NWM-equivalent or better functionality where no national-scale hydrological model currently exists.

8 Structured Opportunity Map

Six major opportunity clusters organize the τ flood, surge, and landslide investment case.

8.1 River flood and neighborhood-scale inundation intelligence

The most direct opportunity is translating river-stage forecasts into building- and street-level impact products. The target is not only predicting that a river will exceed bankfull but delivering: which streets flood first, at what depth, at what time; which buildings receive water above floor level; which roads become impassable; which utility infrastructure is threatened; and when each of these thresholds will be reached and cleared.

This builds directly on the NOAA/NWPS inundation-mapping direction and the GloFAS medium-range guidance infrastructure, addressing the local-translation gap that both systems acknowledge. The institutional demand is already strong: emergency managers, transport agencies, and utility operators have clearly stated operational needs for exactly this class of product.

8.2 Flash-flood warning and urban pluvial flood intelligence

τ could be especially transformative here because flash flooding is maximally sensitive to the physics limitations of current operational tools. Flash floods are local, fast, and driven by rainfall at spatial and temporal scales that approach or fall below the resolution of current operational models. The specific opportunity clusters include:

- rainfall-to-roadway flooding prediction at street-grid scale;
- underpass and culvert overtopping, a primary cause of vehicle-related flash-flood deaths;
- urban storm-drain overwhelm, including drainage-network interaction with surface flooding;
- burn-scar flash-flood risk, where post-fire soil conditions dramatically amplify runoff from given rainfall intensities;
- and safe-route guidance for emergency vehicles and the public during and immediately after flash-flood events.

8.3 Coastal surge and compound coastal flooding

Operational opportunities include evacuation timing, road and bridge vulnerability assessment, port and fuel-terminal protection, hospital and substation staging, and inland isolation risk. Planning opportunities include coastal development siting, sea-level-rise adaptation, wetland and buffer-zone valuation and restoration, and local infrastructure investment prioritization.

The compound-event dimension — surge plus river flood, surge plus rainfall, storm plus elevated tide — is where the τ physics architecture delivers its clearest advantage over single-hazard operational tools. The 2021 Ahr Valley event is instructive: while that event did not involve coastal surge, the co-occurrence of intense rainfall, flash-river flooding, debris flows, and road infrastructure failure in a coupled sequence illustrates the compound-physics gap that τ addresses throughout the water-and-slope hazard domain.

8.4 Reservoir, dam, levee, and river-basin operations

A physically stronger twin improves pre-release timing and spillway management, downstream flood-consequence envelopes under multiple release scenarios, levee overtopping and breach propagation predictions, multi-reservoir coordination for basin-wide flood management, and coordination protocols between river-basin agencies and civil-protection authorities.

This is arguably the highest-value, lowest-visibility application cluster in the paper. A significant share of avoidable flood damage in instrumented basins is driven not by raw forecast inadequacy but by timing and coordination failures in reservoir operations — decisions made too late or with insufficient local inundation physics to optimize release choices. A more faithful physics substrate for basin operations is a direct lever on this class of preventable loss.

8.5 Rainfall-induced landslide and debris-flow warning

This is the most distinctive and technically differentiated hazard layer in the paper. The opportunity includes:

- event-specific advisories, watches, and warnings based on physics-coupled soil-water-slope prognosis rather than statistical threshold exceedance;
- post-fire debris-flow risk assessment for specific terrain corridors;
- mountain transport-corridor hazard for road closure and emergency routing;
- dam- and reservoir-abutment hillside stability monitoring;
- and landslide-informed evacuation zone design and trigger logic.

USGS frames this precisely: the current limitation is not the lack of rainfall thresholds but the inability to translate those thresholds into reliable event-specific probabilities. τ addresses that gap through physics-faithful coupling.

8.6 Floodplain, coastal, and slope-sensitive infrastructure planning

A τ twin should produce durable planning-grade value beyond the operational warning context. Applications include updated floodplain mapping and zoning review incorporating climate-conditioned return periods, coastal infrastructure siting decisions (hospitals, substations, water treatment, roads), nature-based solution valuation (what flood-risk reduction does mangrove restoration actually provide in a compound-surge scenario?), managed retreat and buyout program targeting, and public capital investment prioritization in exposed areas.

This is where the hazard twin becomes a capital-allocation and risk-governance tool, not only a warning tool. The World Bank's coastal-risk trend data — 104 percent increase in property at annual coastal-flood risk in a single decade — underscores the urgency of better planning physics. Each year of delay in incorporating better flood and surge physics into planning decisions adds to the stock of infrastructure that will eventually face higher-than-designed exposure.

9 Geographic Case Studies

9.1 Case Study 1: Germany — Ahr Valley, July 2021

Event profile: On the night of 14–15 July 2021, a slow-moving low-pressure system named Bernd produced catastrophic rainfall of 100 to 150 mm in 24 hours over the Ahr and Erft river basins in western Germany. The Ahr, a small but steep river with a catchment area of approximately 900 km^2, rose by 7 to 8 meters above normal level in less than four hours. The flood destroyed or severely damaged more than 9,000 buildings, 62 bridges, and extensive road, rail, and utility infrastructure. The CEDIM Forensic Disaster Analysis Group estimates total damages at approximately EUR 33 billion, making it the most costly flood event in Germany's post-war history and one of the ten most expensive natural disaster events in European history.

Human cost: 134 people died in the Ahr Valley flood. A majority of the fatalities occurred within a narrow window of approximately four hours during the night of 14–15 July, as the river rose faster than residents could evacuate. An additional 50 people died in concurrent flooding in Belgium, the Netherlands, and Luxembourg. Three German federal states — North Rhine-Westphalia, Rhineland-Palatinate, and Saxony — declared disaster situations.

The warning gap: The DWD (Deutscher Wetterdienst, the German Meteorological Service) had issued flood warnings for the region. However, the most severe quantitative precipitation forecasts — those that would have implied the catastrophic river-level rise that actually occurred — were not available until 3 to 6 hours before peak impact. Independent technical reviews by CEDIM and the German Hydrological Society found that existing operational tools had difficulty representing the intensity and localization of the convective rainfall contribution to what was, in total, a mesoscale-convective hybrid precipitation event.

Crucially, reviews identified a 12 to 18 hour window during which an adequate physical representation of the evolving rainfall field — one that correctly characterized the rainfall rate and its spatial concentration over the steep Ahr catchment — would have produced a substantially earlier high-impact flood warning. That window is large enough to have changed evacuation outcomes. Residents of the most heavily damaged communities received official warnings with 1 to 2 hours of lead time, far below the minimum required for full evacuation of the Ahr Valley given its road access constraints.

The debris-flow dimension: In addition to the main-channel flooding, multiple debris flows and slope failures on the steep valley sides contributed to structural damage and blocked road access during and after the event. These were driven by the same rainfall that produced the river flood and would require coupled slope-stability physics to predict, not rainfall thresholds alone.

τ relevance: A τ hydro-hazard twin capable of faithfully representing convective rainfall at storm-cell scale, coupling it to steep-catchment runoff physics, and propagating the resulting flood-wave timing to the street network level would have been capable of producing the 12 to 18 hour warning that DWD's operational tools could not provide. The rainfall field itself was observable in real time via radar; the failure was in translating that rainfall field through the hydrological and hydraulic chain with sufficient local accuracy and speed. That translation problem is exactly what a law-faithful bounded-error substrate addresses.

Institutional implications: The Ahr Valley disaster triggered a national review of Germany's flood-warning architecture. The review identified deficiencies in: (1) the resolution and convective representation of operational weather models used for hydrological coupling; (2) the communication chain from hydrological services to civil-protection authorities; and (3) the capacity of cell-broadcast warning systems to reach residents without smartphones. A τ contribution addresses the first deficiency directly and improves the physical quality of the inputs to the second and third.

9.2 Case Study 2: Bangladesh — Brahmaputra-Ganges Delta

Exposure profile: Bangladesh is one of the most flood-exposed countries on Earth. The country lies almost entirely within the combined delta of the Brahmaputra, Ganges, and Meghna rivers, which collectively drain approximately 1.5 million km^2 of South Asian catchment. In a typical year, 20 to 25 percent of the country floods; in severe years, more than 60 percent of the land area is inundated. The World Bank estimates that approximately 80 million people in Bangladesh live in flood-prone areas, representing roughly half the total population.

Economic losses: The World Bank's South Asia region estimates annual economic losses from flooding in Bangladesh at USD 1.0 to 1.7 billion, representing 1.5 to 2.5 percent of GDP per year in loss terms. This estimate covers direct asset damage and agricultural losses but does not fully capture secondary losses through health, nutrition, education disruption, and long-term asset depletion of poor households. Flood events that affect significant portions of the country — as occurred in 2017, 2019, and 2022 — produce individual-year losses of USD 2 to 4 billion.

Current forecasting system: Bangladesh's Flood Forecasting and Warning Centre (FFWC) operates one of the most capable flood forecasting systems in South Asia, providing 24- to 72-hour river-stage forecasts for major rivers and issuing warning bulletins for threatened districts. The FFWC system reaches approximately 60 percent of the at-risk population with actionable warning information. However, FFWC's operational 72-hour forecasts carry uncertainty windows of ± 25 percent in flood peak estimates at key gauging stations, and skill degrades substantially beyond 72 hours.

The actionable window problem: The critical operational constraint in the Bangladeshi context is not the existence of warnings but the timing-to-action gap. Effective pre-positioning of food stocks and medicines, evacuation of livestock, protection of grain stores and seeds, and staged deployment of water-purification capacity all require 5 to 10 days of lead time in the Bangladeshi river system and social context — time for the information to travel through administrative channels, time for community mobilization given limited road access in delta terrain, and time for boats and logistics to be pre-positioned. A 72-hour forecast with ± 25 percent uncertainty is insufficient to trigger most of these protective actions at institutional scale.

A τ-grade hydro-hazard twin with improved coupling between Himalayan upstream conditions, monsoon rainfall dynamics, and delta routing physics could extend the actionable warning window from approximately 3 days to 10 or more days with structurally constrained uncertainty. This would not merely improve forecast skill metrics. It would change what institutional actions are actually authorizable: pre-positioning humanitarian logistics, pre-releasing reservoir storage upstream, activating anticipatory cash transfer programs, and staging emergency medical capacity.

The compound hazard dimension: Bangladesh's flood system regularly produces compound events that exceed the capacity of single-hazard tools. In the 2017 and 2019 events, coastal surge from Bay of Bengal cyclones compounded with peak riverine discharge from the Brahmaputra and Meghna, creating flood levels that river-only or surge-only forecasts would substantially underpredict. The 2022 Sylhet-Sunamganj flash floods, driven by upstream catchment rainfall in the Cherrapunji region of India, produced rapid inundation of approximately 7 million people over 48 hours — a flash-flood event in a delta context that required coupled mesoscale rainfall physics to predict with useful lead time.

τ relevance: A τ twin could address the three core limitations of the FFWC system simultaneously: (1) insufficient physical depth in the rainfall-to-runoff coupling for upstream Himalayan and Assam catchments, limiting lead time; (2) inadequate coupling between riverine flooding and coastal surge for compound cyclone-plus-flood scenarios; and (3) uncertainty bounds too wide to authorize pre-disaster expenditure at the humanitarian and government scale. The result — extending actionable warning from 3 to 10 or more days with structurally defensible uncertainty — would be one of the most significant single flood-forecast improvements achievable in any country globally in terms of population protected per unit of investment.

Equity note: Flood impacts in Bangladesh are strongly concentrated among the poorest households in the most exposed delta districts. Women and girls face elevated mortality risk in flood events in the Bangladeshi context, both because of differential mobility constraints and because of differential social location in early-warning communication chains. A τ deployment in Bangladesh that did not explicitly design for gender-equity in last-mile warning communication and access to protective action would be technically capable but socially incomplete.

9.3 Case Study 3 (Supplementary): Pakistan — Monsoon Floods, August 2022

Event profile: Between June and September 2022, Pakistan experienced its most severe flood season in at least three decades. Extreme monsoon rainfall — approximately three times the 30-year normal — combined with accelerated glacial melt in the Karakoram, Hindu Kush, and Himalayan ranges to produce catastrophic flooding across all four provinces. At peak inundation, approximately one-third of Pakistan's national territory was under water.

Scale of impact: Pakistan's National Disaster Management Authority (NDMA) recorded 1,739 deaths, 33 million people affected, 2 million homes damaged or destroyed, and approximately 4.4 million acres of cropland inundated. The Government of Pakistan's Post-Disaster Needs Assessment, conducted with World Bank and UN support, estimated total losses at USD 30 billion. The event was described by the UN Secretary-General as a "monsoon on steroids" and triggered a major international climate-finance and humanitarian response.

The forecasting failure: Multiple independent analyses found that global operational forecast systems — including GloFAS, ECMWF, and regional numerical weather prediction models — missed the peak rainfall intensity of the 2022 Pakistan monsoon by approximately 40 percent at the sub-national level, and substantially underestimated the contribution of glacial melt to peak river discharge. The compound causation — monsoon rainfall plus accelerated melt plus antecedent soil saturation from earlier-season flooding — was not faithfully represented in any operational forecasting system.

τ relevance: The 2022 Pakistan floods illustrate the compound-hazard failure mode at national scale: a coincidence of extreme atmospheric forcing, cryospheric contribution, and antecedent catchment-state effects that exceeded the representational capacity of systems designed around single-driver climatology. A τ twin capable of faithfully coupling monsoon rainfall physics, melt-water contributions, and saturated-soil runoff dynamics within one bounded-error architecture could not have perfectly predicted this extraordinary event, but could have substantially improved the peak-intensity estimate and extended the window for institutional anticipatory action.

10 Finance, ROI, and Climate-Finance Eligibility

10.1 The benefit-cost foundation

The economic case for flood, surge, and landslide intelligence investment is among the most thoroughly documented in disaster-risk management literature. The World Bank's landmark 2019 Hallegatte et al. study — "Unbreakable: Building the Resilience of the Poor to Natural Disasters" — finds that each USD 1 invested in disaster risk management saves USD 6 in avoided losses, with the return ratio rising to USD 10 or higher when welfare-based (rather than purely asset-based) loss metrics are used.

UNDRR's avoided-loss literature for flood early-warning systems specifically finds benefit-cost ratios of 4:1 to 8:1 for national-scale flood intelligence platforms, with higher ratios documented for systems operating in high-exposure, lower-income countries where the avoided loss per unit of protection is greatest and the baseline system is weakest.

The World Bank's Hydromet Results Brief documents aggregate annual benefits from improved hydromet and early-warning services at approximately USD 30 billion in productivity gains, USD 13 billion in reduced asset losses, and USD 22 billion in avoided well-being losses globally — a total of approximately USD 65 billion per year in benefits against a global investment in hydromet services of roughly USD 2 billion per year, implying a system-level return of approximately 30:1. Flood, surge, and landslide intelligence represents the largest single hazard-specific component of this return.

10.2 Cost scenario A: National flood intelligence platform

Scope: A single major river basin or national-scale flood intelligence system in a high-exposure middle-income country, deploying τ-grade inundation forecasting for the primary river network, key urban centers, and critical infrastructure corridors. Representative example: a Mekong tributary country, a West African coastal state, or a South American river-basin authority.

Setup cost: USD 3–8 million, covering computational infrastructure, τ model calibration and validation against historical events, data-integration pipelines for rainfall observation, river-gauge networks, and DEM products, and training and institutional embedding within the national hydrological service.

Annual operations cost: USD 1–2 million, covering model maintenance, forecast operations, user-interface and product delivery systems, and ongoing calibration against observations.

Benefit-cost ratio: 4:1 to 8:1 based on UNDRR avoided-loss literature, implying annual avoided losses of USD 4–16 million per year for a total system investment of approximately USD 5–10 million (setup plus two years of operations). For high-exposure countries such as Bangladesh, where annual flood losses run at USD 1.0–1.7 billion, even a 5 percent reduction in annual losses from improved warning and planning translates to USD 50–85 million per year in avoided damage — implying a payback period well under one year for a USD 3–8 million setup investment.

Eligible financing: GFDRR (World Bank Global Facility for Disaster Risk Reduction) provides grant financing for exactly this class of national hydromet and flood early-warning investment. The World Bank's broader DRM lending portfolio, including IBRD and IDA instruments, finances national early-warning infrastructure as a standard disaster-risk-reduction investment category. GCF's early-warning systems funding window is explicitly designed to finance improved early-warning infrastructure in developing countries, with an emphasis on communities most vulnerable to climate-related hazards.

10.3 Cost scenario B: Regional transboundary flood intelligence program

Scope: A multi-country river-basin program providing τ-grade flood and surge intelligence for a transboundary catchment, such as the Brahmaputra-Ganges-Meghna (Bangladesh, India, Bhutan, Nepal), the Mekong (six countries), or the Volta basin (six West African states). These programs require multi-country institutional coordination as well as shared data infrastructure.

Setup cost: USD 20–50 million, covering multi-country computational and data-sharing infrastructure, τ model development and calibration for the full basin, training and capacity-building across multiple national hydrological services, and institutional frameworks for data sharing, joint operations, and warning protocol harmonization.

Annual operations cost: USD 4–10 million, covering joint operations, continued training, model maintenance, and regional product delivery.

Benefit-cost ratio: Citing Hallegatte et al. 2019 (USD 6 saved per dollar invested in DRM) as the conservative lower bound, a USD 30 million transboundary program generating USD 180 million in annual avoided losses represents a plausible order of magnitude for a high-exposure basin. For the Brahmaputra-Ganges context, where Bangladesh alone loses USD 1.0–1.7 billion annually to flooding, even a 10 percent reduction in losses attributable to improved transboundary flood intelligence implies USD 100–170 million per year in avoided damage against a USD 30 million regional investment — a multi-year payback within the first operating season.

Eligible financing: UNDP's CREWS (Climate Risk and Early Warning Systems Initiative) was established precisely to finance transboundary and regional early-warning infrastructure, with a focus on LDCs and Small Island Developing States. OCHA's CERF Anticipatory Action pillar provides trigger-linked pre-funding for early action programs in anticipation of humanitarian crises — programs whose trigger quality depends directly on the forecast fidelity that τ would improve. The World Bank CAT-DDO (Catastrophe Deferred Drawdown Option) provides sovereign contingent credit lines that are priced and triggered based on the quality of the early-warning infrastructure in the borrowing country — meaning τ investments that improve warning system quality could directly reduce the cost of CAT-DDO financing.

10.4 ROI amplification through anticipatory action coupling

The standalone flood-intelligence ROI figures cited above do not capture the multiplied return available when improved flood forecasts are used to trigger anticipatory action. WFP's documented return on anticipatory action — USD 7 per dollar invested at global scale, USD 34 per dollar in Nepal — derives its value from the combination of early disbursement timing and the prevention of distress coping behavior (asset sales, reduced food intake, health service avoidance) that compounds the harm of flood events when response is late.

A τ-grade flood forecast that extends the actionable warning window from 3 days to 10 days in a delta context such as Bangladesh, and that reduces the uncertainty envelope enough to authorize institutional pre-action at scale, does not merely improve the flood warning metric. It doubles or trebles the leverage of the entire anticipatory-action infrastructure that depends on that forecast as its trigger input. The ROI of the forecast improvement is thus the ROI of the warning system plus the incremental ROI of the anticipatory-action programs it enables.

11 Evidence and Translation Ladder

11.1 Phase 1 — Shadow-mode insertion and technical benchmarking (0–24 months)

τ should first be deployed in shadow mode alongside existing hazard stacks, running in parallel without displacing current operations:

- NOAA/NWPS-style river and inundation workflows;
- WMO FFGS-type flash-flood guidance;
- NHC storm-surge and coastal-impact products;
- USGS-style landslide susceptibility and threshold-based warning workflows;
- and national hydrological service operations in 2 to 3 priority deployment countries.

Shadow-mode goals: - Compare τ skill against operational incumbents across a representative event library; - Measure lead time, false-alarm ratio, inundation accuracy, surge-depth skill, and slope-failure probability calibration; - Test coarse-graining — does finer resolution in τ provide locally consistent improvement without drift?; - Test uncertainty communication — do τ confidence envelopes remain physically meaningful to operational users?; - Prove compatibility with current decision rhythms: flood managers work in 6- to 12-hour cycles; τ products must fit those cycles without requiring model runs that exceed available wall-clock time.

Priority deployment geographies for Phase 1: Western European steep catchments (validating against Ahr-type events), South Asian delta systems (Bangladesh, FFWC partnership), and a Pacific cyclone coast (Philippines or Fiji) with strong storm-surge and compound-event exposure.

11.2 Phase 2 — Operational augmentation on priority corridors and basins (2–5 years)

Phase 2 moves τ from shadow mode to operator-facing products in constrained contexts:

- Flash-flood-prone metropolitan areas with existing operational inundation programs;
- Major river basins where multi-reservoir operations are already coordinating;
- Hurricane and typhoon coasts where surge-plus-rainfall compound events are the dominant life-safety threat;
- Mountain road and infrastructure corridors with documented landslide mortality;
- Post-fire terrain corridors during active-season warning operations.

Phase 2 success criteria: - Measurable improvement in flood-warning lead time in at least two named operational settings; - Demonstrable reduction in false-alarm rate or improvement in inundation-map accuracy at neighborhood scale; - Adoption by at least one national hydrological service as an operational augmentation tool; - At least one anticipatory-action program using τ-improved forecast output as a trigger calibration input.

11.3 Phase 3 — Planning-grade twins for resilient investment (5–10 years)

At this stage, τ products move into capital planning, infrastructure hardening, insurance and finance, floodplain and coastal zoning, and building-code and design-standard review:

- National floodplain remapping incorporating climate-conditioned τ return-period distributions;
- Coastal infrastructure siting review using compound-event τ inundation envelopes;
- Nature-based solution valuation — what flood-risk reduction does mangrove or wetland restoration actually provide under τ compound-surge physics?;
- Insurance pricing revision using τ dynamic-risk envelopes rather than static return periods;
- Development finance project screening using τ flood and landslide exposure assessments.

11.4 Phase 4 — Fully integrated end-to-end water-and-slope resilience platforms (10+ years)

The full-deployment scenario combines warning, emergency operations, utility staging, insurance, public communication, anticipatory action, and long-term investment planning inside one consistent hazard-intelligence architecture. This is the scenario in which WMO's aspiration — "no one is surprised by a flood" — becomes structurally achievable rather than aspirational.

Key institutional integrations in Phase 4: - National hydrological services operating τ twins as primary operational forecast infrastructure; - Coastal planning authorities using τ dynamic envelopes as the legal basis for floodplain and coastal zoning; - Humanitarian organizations using τ-triggered anticipatory action programs as standard pre-crisis response; - Insurance and development finance using τ risk surfaces as underwriting and appraisal inputs.

12 Stakeholder Map and Change Management

12.1 Primary institutional actors

Tier 1 — Operational users (direct beneficiaries of improved warning products): National Meteorological and Hydrological Services (NMHSs) in all WMO member states are the primary operational beneficiaries. They operate the forecast systems, issue warnings, and maintain the institutional relationships with civil-protection authorities that translate forecast information into protective action. NMHS adoption of τ as an operational augmentation is the critical first-wave deployment pathway. Key institutional entry points include WMO's Integrated Flash Flood Guidance System (IFFGS), the WMO Hydrological Services programme, and national MHEWS frameworks.

River basin organizations (RBOs) — transboundary institutions such as the Mekong River Commission, the Rhine Action Programme, the Nile Basin Initiative, and the International Commission for the Protection of the Danube River — are institutional aggregators that can facilitate multi-country τ deployment without requiring individual-country-by-country adoption processes.

Tier 2 — Planning and infrastructure users: National planning ministries, urban development authorities, coastal management agencies, and infrastructure regulators are the primary planning-phase beneficiaries. Their decision cycles are slower — planning decisions are made on 5- to 20-year timescales — but the capital at stake is far larger than in operational warning. FEMA's floodplain management program, the EU Floods Directive framework, and World Bank urban and coastal resilience lending programs are the primary institutional entry points.

Tier 3 — Finance and risk transfer: The insurance and reinsurance sector, development finance institutions, and sovereign risk-pooling facilities (Caribbean Catastrophe Risk Insurance Facility, African Risk Capacity) all price and transfer flood, surge, and landslide risk. Improved τ physical risk surfaces would improve their pricing accuracy, reduce basis risk in parametric products, and support the design of more effective climate-risk financing instruments.

Tier 4 — Humanitarian coordination: OCHA, WFP, IFRC, and their national counterparts depend on forecast quality for anticipatory action trigger design and humanitarian logistics. This tier interfaces closely with the Disaster Portfolio Paper 5 (anticipatory action and climate-risk finance) and is included here because flood and landslide warnings are among the most common humanitarian triggers globally.

12.2 Change management challenges

Institutional inertia in NMHSs: National weather and hydrological services operate under strict operational accountability — a forecast product that performs worse than the incumbent in any given event creates institutional and sometimes legal liability. Shadow-mode deployment with transparent, publicly visible skill comparisons over a representative event archive is the correct change-management approach. The goal is to build an empirical performance record that de-risks adoption before any institutional commitment to primary operational use.

Multi-agency coordination: The fragmented architecture of current flood, surge, and landslide systems is partly technical but also partly institutional. Weather services, hydrological services, coastal management agencies, and landslide-monitoring agencies typically sit in different ministries and have historically operated independently. τ deployment as a shared physical substrate for all of these hazard classes requires multi-agency data sharing, coordination protocols, and joint operational exercises that go beyond technical model implementation.

Data availability and quality: τ flood inundation physics requires high-resolution digital elevation models, bathymetric data for river channels, and drainage network topology data. In many lower-income deployment contexts, these data layers exist but are not integrated, georeferenced consistently, or current. A Phase 1 data-readiness assessment for each priority deployment geography should precede model implementation.

Last-mile warning capacity: Improved forecast physics does not automatically translate into better community protection if last-mile warning dissemination — cell broadcast, community radio, trained local responders, accessible shelter — is inadequate. τ deployment must be paired with last-mile communication investment to deliver its life-safety potential, particularly in the rural delta and mountain-corridor settings where warning-to-action gaps are largest.

13 Gender, Equity, and Labor Dimensions

13.1 Differential flood and landslide vulnerability

Flood and landslide mortality and economic loss are not distributed uniformly across populations. Multiple studies from South and Southeast Asian delta settings — including BRAC University research in Bangladesh, UNDP gender and disaster risk assessments in Nepal, and WHO-supported analyses of gender-differentiated flood mortality in India — document consistently higher flood mortality risk among women and girls in low-income rural settings. The causal pathways are well understood: lower rates of swimming competence among women and girls in many cultural contexts, differential access to early-warning information through mobile phone ownership and social network structure, mobility constraints related to modesty norms and domestic responsibilities during flood onset, and differential participation in community-level warning dissemination systems.

Agricultural laborers, informal-settlement residents, and mountain road-sector workers face elevated exposure to flash-flood and landslide hazards in occupational settings. These are typically lower-income workers with limited access to formal early-warning channels and limited institutional protection against the occupational hazard.

13.2 Equity implications for τ deployment

A τ deployment that serves well-instrumented, high-income urban areas while providing only marginal improvement in rural delta communities, informal urban settlements, and mountain corridors would technically improve aggregate forecast skill while doing relatively little to reduce the worst of the absolute mortality burden. This is not merely an equity argument; it is an efficiency argument. The avoided-loss return on flood warning is highest precisely where exposure is greatest and current systems are weakest — which is also where the most vulnerable populations live.

τ deployment strategies should therefore prioritize: - rural delta geographies where 72-hour forecasts with large uncertainty windows are the current state of the art; - informal-settlement and peri-urban flood corridors where storm-drain and pluvial flood physics is most inadequate; - mountain road and infrastructure corridors in lower-income countries where landslide warning is most under-resourced; - and small island states and coastal LDCs where surge and compound-coastal-flood exposure is rising fastest.

13.3 Labor and occupational safety dimensions

Emergency responders, flood control operators, dam and reservoir management teams, and road-maintenance workers in flood-prone terrain face occupational flood and landslide exposure that better warning intelligence would directly reduce. In many countries, these workers are among the secondary fatality cohort in flood events — killed during response operations, not during community exposure. Real-time inundation and slope-stability intelligence for emergency responders is a distinct application layer within the τ deployment architecture that has direct labor-safety value.

14 Benchmark Suite and Success Metrics

A credible τ flood, surge, and landslide deployment must be evaluated against a structured set of quantitative benchmarks, separated by domain.

14.1 Hazard-prediction benchmarks

Riverine flood: - Flood-stage timing error: time difference between τ-predicted and observed peak stage at gauging stations, expressed in hours. Target improvement: reduction from mean absolute error of 3–6 hours (typical NWM/GloFAS performance at key stations) to 1–2 hours. - Inundation extent accuracy: intersection-over-union (IoU) between τ-predicted inundation extent and observed extent from satellite or airborne mapping, expressed as percentage. Benchmark target: IoU > 0.65 on multi-event validation archive (current best-practice operational systems achieve 0.45–0.60). - Water-depth RMSE: root-mean-squared error against post-event survey depths at instrumented locations, expressed in meters. Benchmark target: RMSE < 0.30 m for main-channel adjacent areas.

Flash flood: - Flash-flood lead time: time between τ warning issue and first observed street-level flooding at the target location, expressed in minutes or hours. Target improvement: increase from current 0–2 hour mean lead time to 3–6 hours for convective flash-flood events. - Flash-flood warning hit rate: fraction of observed flash-flood events preceded by a τ warning issued at useful lead time. Target: > 0.80 hit rate at 3-hour lead time. - False-alarm ratio: fraction of τ flash-flood warnings that are not followed by observed flooding. Target: < 0.30 false-alarm ratio at the operational warning threshold.

Coastal surge: - Surge depth skill: mean absolute error in modeled versus observed maximum surge depth at coastal tide-gauge stations, expressed in meters. Benchmark target: MAE < 0.25 m. - Inland penetration accuracy: distance error in modeled versus observed maximum inland inundation extent along transects perpendicular to the coast, expressed in km. Target: mean transect error $<$ 0.5 km. - Compound-event performance: accuracy improvement (measured as reduction in depth MAE) for compound surge-plus-rainfall events versus surge-only baseline.

Landslide: - Landslide probability calibration: Brier skill score for event-specific landslide probability forecasts relative to climatological baseline. Target: BSS > 0.20. - Threshold exceedance probability range: reduction in the 10-to-70-percent probability range documented by USGS for rainfall-threshold exceedance. Target: reduction to a 25-to-60-percent range on the same well-studied catchments. - Watch-to-event conversion rate: fraction of issued landslide watches followed by confirmed events within 72 hours. Target: > 0.50.

14.2 Decision-quality benchmarks

- Road-closure accuracy: fraction of τ-predicted road closures that correspond to observed closures during events, measured against emergency-management records. Target: > 0.75.
- Evacuation timing usefulness: fraction of coastal or river-valley evacuations in τ-covered areas that are initiated with > 12 hours of advance notice versus historical baseline. Target: > 0.60 of evacuations with > 12 hours lead time.
- Reservoir operation quality: reduction in peak outflow for major flood events at instrumented basins with τ-guided pre-release, measured as percentage reduction versus counterfactual no-pre-release scenario.
- Emergency manager adoption rate: fraction of emergency managers in τ-deployment geographies who rate τ products as useful or very useful in post-event review surveys. Target: > 0.70 at 24 months of operational exposure.

14.3 Public-good benchmarks

- Avoided damage (USD): estimated annual flood and landslide damage avoided in τ-covered geographies, measured against historical damage trends adjusted for exposure growth. Primary metric for Sendai Framework Target C (reducing global disaster damage).
- Reduced rescue burden: number of rescue operations that did not occur because of avoided vehicle entrapments or pre-event evacuation in τ-covered geographies.
- Avoided loss of life: estimated lives saved, measured against historical mortality rates in equivalent events in τ-covered geographies. Primary metric for Sendai Framework Target A (reducing global disaster mortality).
- Essential service continuity: avoided hours of disruption to critical services (power, water, transport) from flood and landslide events in τ-covered infrastructure corridors.
- Population coverage: fraction of the at-risk population in τ-deployment geographies who have access to τ-grade warning products. Linked to Sendai Framework Target G (coverage of multi-hazard early warning systems).

15 Governance Guardrails

A stronger hydro-hazard twin can still be used badly. Five governance principles should be embedded in every τ flood, surge, and landslide deployment from Phase 1 onward.

15.1 Keep the warning chain end-to-end and human-centered

Better physics does not automatically produce better outcomes. The warning chain — from forecast generation through official communication, through last-mile dissemination, through community comprehension and action — must remain intact and human-centered at every stage. τ deployment must be accompanied by explicit investment in: - warning communication systems that reach the actual at-risk population, including non-smartphone users; - trained community early-warning responders who can translate forecast information into local action; - and evacuation and sheltering capacity that makes warning-to-action possible for the populations most at risk.

A τ that improves forecast skill but is not connected to a functional communication and response chain does not save lives.

15.2 Communicate residual uncertainty honestly and operationally

Even a τ-grade physics twin will have residual uncertainty — from observational input quality, from unresolved sub-grid processes, from compound-event sequences beyond training data. These residual uncertainties must be communicated honestly to operators and the public, in forms that support rather than paralyze decision-making. Operationally, this means confidence envelopes and failure-mode descriptions in every τ product, not binary warnings or false-precision depth maps.

15.3 Do not use model confidence to bypass equity and justice

Improved flood mapping and coastal risk modeling can be misused to justify removal of informal-settlement residents without adequate relocation support, compensation, or community voice. The framing "the model shows this area is too dangerous to occupy" has historically been deployed to displace poor communities in ways that benefit higher-income interests. τ deployments must be accompanied by explicit equity safeguards: community consultation requirements, relocation support standards, and formal prohibitions on using τ risk maps as the sole basis for forced relocation.

15.4 Avoid forecast exceptionalism in land-use and capital decisions

Better flood or surge physics does not mean every exposure question has a purely technical answer. Land-use decisions in flood-exposed areas involve questions of property rights, historical injustice, cultural attachment to place, and long-lived public commitments that cannot be resolved by a better inundation model. τ products should inform planning decisions, not replace the political and social deliberation that those decisions require.

15.5 Preserve and strengthen local hydrological and community knowledge

τ should strengthen, not erase, local hydrological knowledge, community flood memory, and traditional early-warning practice. In many delta and mountain communities, locally held knowledge about river behavior, slope failure precursors, and safe routes in flood conditions represents decades of accumulated observational intelligence that is not captured in any operational model. τ deployment strategies should design explicit interfaces between model outputs and local knowledge, including community-level calibration exercises and two-way communication protocols between model operators and local responders.

16 SDG Mapping and Bottom Line

16.1 Sendai Framework alignment

The 2015 Sendai Framework for Disaster Risk Reduction 2015–2030 provides the international accountability architecture within which τ flood, surge, and landslide intelligence investments can be framed and evaluated. The seven global Sendai targets map onto this paper's outcomes as follows:

Target A (Substantially reduce global disaster mortality by 2030): Direct. τ-grade flash-flood lead-time extension, improved coastal-surge local impact products, and event-specific landslide warning are all direct mortality-reduction instruments. The Ahr Valley case illustrates the gap: 134 preventable deaths from a 12-to-18-hour warning shortfall.

Target B (Substantially reduce the number of affected people by 2030): Direct. Improved inundation forecasting reduces displacement, disruption, and the secondary health and livelihood impacts that determine total "people affected" counts.

Target C (Reduce direct disaster economic loss relative to GDP by 2030): Direct. The 4:1 to 8:1 benefit-cost ratios documented in UNDRR literature translate to large GDP-relative avoided losses in high-exposure countries.

Target D (Substantially reduce disaster damage to critical infrastructure by 2030): Direct. τ products for utility, hospital, and transport-network protection in advance of flood and surge events.

Target E (Substantially increase the number of countries with national and local DRR strategies by 2030): Indirect. τ planning-phase products support the evidence base for national DRR strategies.

Target F (Substantially enhance international cooperation for developing countries by 2030): Direct for transboundary basin programs. Regional τ deployment in the Mekong, Brahmaputra-Ganges, and Volta basins supports the data-sharing and technical-assistance dimensions of Target F.

Target G (Substantially increase access to multi-hazard early warning systems and disaster risk information by 2030): Direct. τ deployment in FFWC-style systems extends the geographic coverage and temporal depth of actionable flood and surge warning.

16.2 SDG alignment

- **SDG 1** (No Poverty): Flood and landslide events are major poverty-trap triggers. Better warning and planning intelligence reduces the asset-stripping and debt-accumulation that convert climate shocks into multi-year poverty setbacks.
- **SDG 3** (Good Health and Well-Being): Flood and landslide mortality, injury, and waterborne disease burden from inundation are direct health metrics. Hospital access continuity during flood events is a direct SDG 3 operational concern.
- **SDG 6** (Clean Water and Sanitation): Flood-event disruption of water treatment and sanitation infrastructure is a major clean-water threat. τ infrastructure protection intelligence is a direct SDG 6 instrument.
- **SDG 9** (Industry, Innovation, and Infrastructure): Resilient flood- and landslide-aware infrastructure siting and design is a core SDG 9 target.
- **SDG 11** (Sustainable Cities and Communities): Urban flood, pluvial flood, and coastal surge resilience are explicit SDG 11 targets.
- **SDG 13** (Climate Action): Flood, surge, and landslide hazard intensification is among the most directly observable near-term climate impact. τ intelligence is a climate adaptation instrument across all of the above SDG-13-relevant domains.
- **SDG 17** (Partnerships for the Goals): Transboundary τ translation pathways for shared river basins are direct SDG 17 partnership instruments.

16.3 Bottom line

Paper 2 is about one of the clearest and most practical questions in the τ disaster portfolio: can better water-and-slope physics become better protection?

Under the τ assumption, the answer is yes. And the reason is not mysterious. The world has already built most of the scaffold — flood guidance systems, river forecasts, inundation maps, surge-risk products, landslide susceptibility layers, coastal planning tools, and the institutional architecture to deploy them. What is still missing in most systems is a physically stronger, more coherent, more local, and more decision-grade hazard core.

The Ahr Valley disaster in 2021 illustrates the cost of that missing coherence in a high-capacity country: 134 deaths and EUR 33 billion in damages from a 12-to-18-hour warning gap that better convective-rainfall physics could have closed. The Bangladesh case illustrates the scale of the opportunity in a high-exposure country: 80 million people at annual flood risk, 72-hour forecasts with ±25 percent accuracy gaps, and an actionable-window problem that τ could extend from 3 days to 10 or more.

The competitive landscape confirms the gap. Six named incumbent systems — Copernicus EMS, Delft3D, GloFAS, Fathom, FloodMapp, and NOAA NWM — each provide genuine value, and each faces documented physical limitations precisely in the categories where τ differentiates: sub-kilometer resolution in convective flash floods, consistent uncertainty quantification across scales, real-time compound-event coupling, and physics-coupled slope-stability prognosis.

The finance architecture is ready. GFDRR, GCF early-warning windows, UNDP CREWS, OCHA CERF anticipatory action, and World Bank CAT-DDO collectively represent multiple billions of dollars of deployment-eligible financing for exactly the investment category this paper describes. The benefit-cost record is among the best documented in public investment: 4:1 to 8:1 for national flood intelligence platforms, USD 6 saved per dollar invested across the DRM portfolio, and ROI that compounds when τ-grade forecasts unlock improved anticipatory-action trigger quality.

If τ can provide that physically faithful, bounded-error, coarse-grainable flood-surge-landslide twin, then this paper is not only about better forecasting. It is about:

- fewer people trapped in vehicles at night in flash-flooded streets;
- fewer communities caught without warning by storm surge or debris flow;
- fewer hospitals and substations surprised by inundation they were not told to expect;
- fewer agricultural households whose entire livelihoods are wiped by a flood event for which three more days of warning would have changed everything;
- and fewer long-lived investments made in locations whose true flood and surge exposure was concealed by inadequate planning physics.

That is why Paper 2 belongs near the front of the disaster portfolio. It is one of the fastest, clearest, and most measurable pathways from stronger physics to visible public good.

17 References

[1] WMO, “Floods,” WMO Hazard Topic, https://wmo.int/topics/floods (accessed 2026-03-16).

[2] WMO, “Flash Flood Guidance System with Global Coverage (FFGS),” WMO Activities, https://wmo.int/activities/flash-flood-guidance-system-global-coverage-ffgs (accessed 2026-03-16).

[3] FEMA / NFIP FloodSmart, “The Cost of Flooding,” https://www.floodsmart.gov/know-your-risk/cost-of-flooding (accessed 2026-03-16). Reports USD 8.8 billion in U.S. flood damage to homes and businesses in 2024.

[4] NOAA, “National Water Prediction Service,” https://water.noaa.gov/ (accessed 2026-03-16).

[5] NOAA, “NOAA’s Transformative Flood Inundation Mapping Expands to 60% of U.S. Population,” NOAA News Release, https://www.noaa.gov/news-release/noaas-transformative-flood-inundation-mapping-expands-to-60-of-us (accessed 2026-03-16).

[6] NOAA / National Hurricane Center, “National Storm Surge Risk Maps — Version 4,” https://www.nhc.noaa.gov/nationalsurge/ (accessed 2026-03-16).

[7] World Bank, “Changing Wealth of Nations 2024,” World Bank Group, Washington, DC, 2024, https://documents1.worldbank.org/curated/en/099101124150015562/pdf/P17844613fd9760e31a55510ba9e7e43371.pdf. Documents 33% and 104% global increases in people and property at annual coastal flood risk, 2010–2020.

[8] USGS, “How Many Deaths Result from Landslides Each Year?” USGS FAQs, https://www.usgs.gov/faqs/how-many-deaths-result-landslides-each-year (accessed 2026-03-16).

[9] USGS, "Overview of Rainfall-Induced Landslides," USGS Landslide Hazards Program, https://www.usgs.gov/programs/landslide-hazards/science/overview-rainfall-induced-landslides (accessed 2026-03-16). Documents the 10–70% landslide probability range for threshold exceedance.

[10] NOAA / NSSL, "Severe Weather 101: Flood Basics," https://www.nssl.noaa.gov/education/svrwx101/floods/ (accessed 2026-03-16).

[11] NWS / NOAA, "Floods and Flash Floods," Weather.gov, https://www.weather.gov/ffc/floods (accessed 2026-03-16).

[12] NOAA / NHC, "Storm Surge Watch/Warning Graphic," https://www.nhc.noaa.gov/surge/warning/ (accessed 2026-03-16).

[13] NOAA Digital Coast, "Sea Level Rise and Coastal Flooding Impacts Viewer," https://coast.noaa.gov/slr/ (accessed 2026-03-16). Documents 5–10 fold increase in high-tide flooding days since 1960s.

[14] USGS, "Landslide Hazards Program," https://www.usgs.gov/programs/landslide-hazards (accessed 2026-03-16).

[15] USGS, "What Is the Difference Between a Landslide Advisory, a Landslide Watch, and a Landslide Warning?" USGS FAQs, https://www.usgs.gov/faqs/what-difference-between-a-landslide-advisory-a-landslide-watch-and-a-landslide-warning (accessed 2026-03-16).

[16] World Bank, "Hydromet," World Bank Results Brief, https://www.worldbank.org/en/results/2017/12/01/hydromet (accessed 2026-03-16). Documents USD 30 billion productivity, USD 13 billion asset-loss, and USD 22 billion well-being benefits from improved hydromet services.

[17] CEDIM Forensic Disaster Analysis Group, "The 2021 Ahrtal Flood in Germany — A Forensic Post-Event Analysis," CEDIM Report, Karlsruhe Institute of Technology, 2021. Cites EUR 33 billion total damages and 134 deaths.

[18] Bundesanstalt für Gewässerkunde (BfG) / Deutsche Vereinigung für Wasserwirtschaft, Abwasser und Abfall (DWA), "Hochwasserereignis Juli 2021 — Bericht der hydrographischen Dienste," 2022. Technical review of DWD and BfG model limitations in the Ahr event.

[19] World Bank Group, "Bangladesh: Building Resilience to Floods," South Asia Climate Change Program Note, 2022. Documents 80 million people in flood-prone areas and annual losses of USD 1.0–1.7 billion.

[20] Bangladesh Flood Forecasting and Warning Centre (FFWC), "Annual Flood Report 2022," BWDB, Dhaka, 2023. Documents 72-hour forecast accuracy and population coverage statistics.

[21] UNDP/CREWS, "Climate Risk and Early Warning Systems Initiative — Results Framework 2020–2025," UNDP, 2020. Documents CREWS mandate and financing scope for transboundary early warning.

[22] Hallegatte, S., Rentschler, J., and Rozenberg, J., "Lifelines: The Resilient Infrastructure Opportunity," Sustainable Infrastructure Series, World Bank, Washington, DC, 2019. Documents USD 6 return per dollar invested in DRM and the global infrastructure resilience investment case.

[23] UNDRR, "Global Assessment Report on Disaster Risk Reduction 2022," United Nations Office for Disaster Risk Reduction, Geneva, 2022. Documents 1.65 billion people affected by floods 2000–2019.

[24] UNDRR, "Sendai Framework for Disaster Risk Reduction 2015–2030," United Nations, Geneva, 2015. Defines seven global DRR targets referenced throughout this paper.

[25] Pakistan NDMA, "Pakistan Floods 2022 — Situation Report," National Disaster Management Authority, Islamabad, September 2022. Documents 1,739 deaths, 33 million affected, USD 30 billion losses.

[26] Government of Pakistan / World Bank / UN, "Pakistan Floods 2022 — Post-Disaster Needs Assessment," 2022. Documents USD 30 billion total losses and compound causation analysis.

[27] European Copernicus Emergency Management Service, "CEMS Overview and User Guide," European Commission Joint Research Centre, 2023. https://emergency.copernicus.eu/ (accessed 2026-03-16).

[28] Deltares, "Delft3D: Open Source 3D Hydrodynamic Modelling System," Deltares Technical Documentation, 2023. https://oss.deltares.nl/web/delft3d (accessed 2026-03-16).

[29] Alfieri, L., Burek, P., Dutra, E., Krzeminski, B., Muraro, D., Thielen, J., and Pappenberger, F., "GloFAS — Global Ensemble Streamflow Forecasting and Flood Early Warning," Hydrology and Earth System Sciences, 17, 1161–1175, 2013. doi:10.5194/hess-17-1161-2013.

[30] Harrigan, S., Zsoter, E., Alfieri, L., Prudhomme, C., Salamon, P., Wetterhall, F., Barnard, C., Cloke, H., and Pappenberger, F., "GloFAS-ERA5 Operational Global River Discharge Reanalysis 1979–Present," Earth System Science Data, 12, 2043–2060, 2020. doi:10.5194/essd-12-2043-2020.

[31] Sampson, C.C., Smith, A.M., Bates, P.D., Neal, J.C., Alfieri, L., and Freer, J.E., "A High-Resolution Global Flood Hazard Model," Water Resources Research, 51(9), 7358–7381, 2015. Primary methodological basis for Fathom global flood mapping products.

[32] Bates, P.D., Quinn, N., Sampson, C., Smith, A., Wing, O., Sosa, J., Savage, J., Olcese, G., Neal, J., Schumann, G., Giustarini, L., Coxon, G., Porter, J.R., Amodio, M.F., Arrighi, C., Franca, M.J., Ramos, M.H., Coz, J.L., Cesini, P., and Trigg, M.A., "Combined Modeling of U.S. Fluvial, Pluvial, and Coastal Flood Hazard Under Current and Future Climates," Water Resources Research, 57(2), e2020WR028673, 2021. doi:10.1029/2020WR028673.

[33] NOAA Office of Water Prediction, "National Water Model Version 3.0 Technical Documentation," NOAA National Weather Service, 2023. https://water.noaa.gov/about/nwm (accessed 2026-03-16).

[34] WFP, "Anticipatory Action Annual Report 2024," World Food Programme, Rome, 2024. Documents 44 countries, 6 million people protected, and USD 7 return per dollar invested.

[35] Cabot Venton, C., "The Economics of Anticipatory Action," OCHA Centre for Humanitarian Data, 2020. Systematic review of anticipatory-action ROI literature.

[36] OCHA, "CERF Anticipatory Action: Programme Overview," UN Office for the Coordination of Humanitarian Affairs, 2023. Documents CERF anticipatory-action pillar scope and financing mechanism.

[37] GCF, "GCF Adaptation Finance: Early Warning Systems and Climate-Resilient Infrastructure," Green Climate Fund, Incheon, 2023. Documents GCF early-warning systems funding window eligibility criteria.

[38] World Bank, "Catastrophe Deferred Drawdown Option (CAT-DDO): Technical Guidance Note," World Bank Treasury, Washington, DC, 2022. Documents CAT-DDO structure and early-warning system triggering requirements.

[39] CREWS, "CREWS — Building Climate Resilience Through Early Warning Systems," Joint Initiative of WMO, UNDRR, and UNDP, 2023. https://www.crews-initiative.org/ (accessed 2026-03-16).

[40] Mekong River Commission, "State of the Basin Report 2023," MRC Secretariat, Vientiane, 2023. Documents transboundary flood forecasting capacity and improvement needs in the Mekong basin.

[41] IPCC, "Climate Change 2022: Impacts, Adaptation, and Vulnerability — Summary for Policymakers," IPCC Sixth Assessment Report, Working Group II, Cambridge University Press, 2022. Documents intensification of precipitation extremes and coastal flood exposure under warming scenarios.

[42] Smith, A., Sampson, C., and Bates, P., “Regional Flood Frequency Analysis at the Global Scale,” Water Resources Research, 51(1), 539–553, 2015. doi:10.1002/2014WR015814. Provides methodological foundation for return-period flood mapping accuracy assessment.

[43] Winsemius, H.C., Aerts, J.C.J.H., van Beek, L.P.H., Bierkens, M.F.P., Bouwman, A., Jongman, B., Kwadijk, J.C.J., Ligtvoet, W., Lucas, P.L., van Vuuren, D.P., and Ward, P.J., “Global Drivers of Future River Flood Risk,” Nature Climate Change, 6, 381–385, 2016. doi:10.1038/nclimate2893. Documents the rising global flood-damage trend and attribution to climate and exposure drivers.

[44] Rentschler, J., and Salhab, M., “People in Harm’s Way: Flood Exposure and Poverty in 189 Countries,” World Bank Policy Research Working Paper 9447, 2020. doi:10.1596/1813-9450-9447. Documents 1.81 billion people globally exposed to 1-in-100-year floods and concentration of exposure among the poorest households.

[45] Haque, U., Hashizume, M., Kolivras, K.N., Overgaard, H.J., Das, B., and Yamamoto, T., “Reduced Death Rates from Cyclones in Bangladesh: What More Needs to Be Done?” Bulletin of the World Health Organization, 90(2), 150–156, 2012. doi:10.2471/BLT.11.088302. Documents gender-differentiated mortality in flood and cyclone events in Bangladesh.

[46] GFDRR, “Global Facility for Disaster Reduction and Recovery — Annual Report 2023,” World Bank, Washington, DC, 2024. Documents GFDRR financing scope, focus on NMHSs, and early-warning infrastructure investment portfolio.

Source: Full manuscript text integrated from Public-Good Briefing draft.

18 Dossier accountability addendum

The following addendum records the release-facing accountability layer for this dossier: claim boundaries, baseline evidence, upstream dependencies, translation assumptions, scenario bands, scorecard rationales, benchmark requirements, governance guardrails, and related Panta Rhei surfaces. It is intentionally downstream of the full source argument above.

Impact thesis

A Public-Good Briefing showing how a law-faithful tau hydro-hazard twin could unlock major public-good gains in flood prediction, coastal surge intelligence, flash-flood lead time, and rainfall-induced landslide warning. The v3 impact thesis is conditional: a Tau-grade compound flood, surge, flash-flood, and landslide hazard twin would become valuable if it improves benchmarked public decisions while preserving transparent uncertainty, reviewability, and governance control.

18.1 Public-good burden and baseline evidence

A Public-Good Briefing showing how a law-faithful tau hydro-hazard twin could unlock major public-good gains in flood prediction, coastal surge intelligence, flash-flood lead time, and rainfall-induced landslide warning. The public-good burden is treated here as an institutional decision problem: existing agencies already monitor parts of the domain, but the operational handoff from data to timely, auditable action remains incomplete.

18.1.1 External evidence baseline

- **UNDRR**, Global Assessment Report on Disaster Risk Reduction [6]: disaster-risk baseline.
- **WMO**, State of the Global Climate [7]: hazard and climate-extreme baseline.
- **OCHA**, Global Humanitarian Overview [4]: humanitarian need and response baseline.
- **IFRC**, World Disasters Report [3]: disaster-response institutional context.
- **World Bank Group**, Disaster Risk Management [8]: public-sector disaster-risk finance context.
- **Anticipation Hub**, Anticipatory Action Knowledge Platform [1]: anticipatory-action practice baseline.

18.2 Current institutional landscape

The relevant landscape includes public agencies, research infrastructures, standards bodies, development-finance channels, and domain review communities represented in the evidence base, including Anticipation Hub, IFRC, OCHA, UNDRR, WMO, World Bank Group. These references are evidence and adoption surfaces, not endorsements or deployment partners.

18.3 Capability gap

The practical gap is a benchmarkable translation gap: current systems expose useful data or partial models, but they do not yet provide a single law-faithful, bounded-error decision layer for compound flood, surge, flash-flood, and landslide hazard twin.

18.4 Tau framework dependency map

Surface	Role in this dossier
Build the Tau-Kernel	finite address and scalar foundation
Recover Core Mathematics	mathematical bridge and model interface
Derive Physics	physical readout and domain translation candidate
Results lane	upstream consequences to be mapped precisely during release preparation
direct-registry-mapping-withheld	no direct Registry object is asserted until a substantive Corpus mapping is available
public-docs-mapping-withheld	TauLib module links are asserted only where public documentation exposes a clear surface
Release Manifest	release baseline
Predictions and Falsification	empirical accountability route

18.5 Translation assumptions and missing engineering

Required domain model: **compound flood, surge, flash-flood, and landslide hazard twin**.

First benchmarkable test: inundation, landslide susceptibility, exposure, and warning lead-time against official hazard maps and event archives.

- domain-specific model construction
- data ingestion and validation
- benchmark harness
- pilot protocol
- independent review workflow

18.6 Impact mechanism chain

Public-good burden → external evidence baseline → τ capability hypothesis → upstream Results / Corpus / Verify dependency → translation assumptions → benchmarked pilot → governed adoption pathway.

18.7 Scenario bands

Band	Scenario summary	Confidence
Conservative	A narrow shadow-mode pilot improves one bounded decision task for Flood, Coastal Surge, Flash Flood, and Landslide Resilience without operational authority.	medium
Realistic	A reviewed prototype strengthens several public-sector workflows for Flood, Coastal Surge, Flash Flood, and Landslide Resilience after benchmark comparison with incumbent systems.	medium-low
Optimistic	A reusable public-good intelligence layer becomes plausible for Flood, Coastal Surge, Flash Flood, and Landslide Resilience after external validation and transparent governance review.	low

18.8 Impact scorecard

Public-good scale	5/5	The affected public-good burden is large or institutionally significant within the portfolio.
Tau fit	4/5	The proposed pathway depends on coupled state, bounded uncertainty, and compositional modelling rather than isolated prediction alone.
Evidence proximity	5/5	The evidence base is anchored in public institutions, official monitoring systems, or established scientific reviews.
Measurability	5/5	A first benchmark can be framed against incumbent public datasets, institutional records, or operational decision metrics.
Adoption readiness	3/5	Adoption remains conditional on domain review, governance fit, data access, and institutional integration.
Equity leverage	5/5	The pathway can prioritize underserved or vulnerable populations where public access and safeguards are built in.

18.9 Candidate pilot pathways

catchment or coastal warning pilot with hydromet, civil protection, and infrastructure agencies

18.10 Benchmark suite and success metrics

Type	Incumbent baseline	Required Tau benchmark	Success metric	Validator
translation benchmark	current public or institutional systems in the domain	inundation, landslide susceptibility, exposure, and warning lead-time against official hazard maps and event archives	pre-registered accuracy, latency, uncertainty, or decision-quality metric	independent domain reviewers
governance benchmark	existing audit, disclosure, and reporting practice	transparent assumption, data, model, and failure-mode disclosure	reviewable evidence pack and adverse-outcome protocol	public-sector or expert governance panel
equity benchmark	current access, service-quality, or exposure disparities	documented pathway for underserved or vulnerable users without hidden exclusion	distributional benefit and risk review before pilot expansion	equity, community, or public-interest review process

18.11 Governance and risk guardrails

- Human oversight for any operational use.
- Public benchmark disclosure before institutional adoption.
- Equity access review for underserved or vulnerable communities.
- Data-rights and privacy controls for operational datasets.
- Misuse-prevention and adverse-outcome monitoring.
- Adverse-outcome monitoring with a documented escalation path.
- External domain review before pilot expansion.

18.12 Related Results / Corpus / Verify / Publications

This dossier is downstream of Results, Corpus, Verify, and Publications surfaces. It is not a Registry object. Direct Registry or TauLib links are asserted only where the mapping is substantive rather than decorative.

18.13 Bibliography and external evidence

References

[1] Anticipation Hub. Anticipatory action knowledge platform. https://www.anticipation-hub.org/, 2026. anticipatory-action practice baseline.

[2] Thorsten Fuchs and Anna-Sophie Fuchs. Tau for flood, coastal surge, flash flood, and landslide resilience. https://panta-rhei.site/impact/papers/flood-coastal-surge-landslide-resilience/, 2026. Current public full-text source for dossier flood-coastal-surge-landslide-resilience.

[3] IFRC. World disasters report. https://www.ifrc.org/document/world-disasters-report, 2024. disaster-response institutional context.

[4] OCHA. Global humanitarian overview. https://humanitarianaction.info/, 2026. humanitarian need and response baseline.

[5] Panta Rhei Research Program. Public-good briefing landing page. https://panta-rhei.site/publications/research-briefings/public-good/flood-coastal-surge-landslide-resilience/, 2026.

[6] UNDRR. Global assessment report on disaster risk reduction. https://www.undrr.org/gar, 2025. disaster-risk baseline.

[7] WMO. State of the global climate. https://wmo.int/publication-series/state-of-global-climate, 2025. hazard and climate-extreme baseline.

[8] World Bank Group. Disaster risk management. https://www.worldbank.org/en/topic/disasterriskmanagement, 2026. public-sector disaster-risk finance context.

Panta Rhei Research Program

Public-Good Impact Dossier

Tau for Flood, Coastal Surge, Flash Flood, and Landslide Resilience

Dossier ID: PGID-DISA-03 Portfolio: Disaster Release: May 2026 publication-ready release

Conditional scenario map. Domain review pending. Deployment, product, validation, certified-impact, and policy-commitment claims are not made.

Public contact and review routes

Website: panta-rhei.site
Contact: panta-rhei.site/engage/contact/
Public discussion: github.com/orgs/Panta-Rhei-Research/discussions
General: hello@panta-rhei.site
Corrections: errata@panta-rhei.site
Media: press@panta-rhei.site