Panta Rhei
Research Program

Research Briefings · Public-Good Impact Dossiers

Agriculture · **Food, Life & Health Systems**

Tau for Seasonal Planning, Disaster Anticipation, and Food-System Resilience

Conditional public-good pathway for Seasonal Planning, Disaster Anticipation, and Food-System Resilience

Public-Good Impact Dossier

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

Dossier ID: PGID-AGRI-04 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-AGRI-04
Portfolio	Agriculture
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/seasonal-planning-disaster-anticipation-food-system-resilience/
Landing route	https://panta-rhei.site/publications/research-briefings/public-good/seasonal-planning-disaster-anticipation-food-system-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.

Contents

1 Executive Summary

This dossier examines one of the highest-stakes applications in the entire τ agriculture portfolio: whether a physically faithful, bounded-error, coarse-grainable seasonal planning twin can close the gap between climate science and food-system protection at scale. The answer, under the explicit τ assumption, is affirmative — and the public-good scale is substantial.

The institutional need is already documented. FAO's 2025 disaster assessment records USD 3.26 trillion in agricultural losses from 1991–2023, averaging USD 99 billion per year, with nearly USD 2.9 trillion attributable to climate-related hazards. WFP's anticipatory-action portfolio now spans 44 countries and protected over 6 million people in 2024, delivering up to USD 7 in avoided losses for every USD 1 invested. The Global Report on Food Crises 2025 reports that more than 295 million people across 53 countries faced acute hunger in 2024. The World Bank now characterizes the dominant failure mode in food-system resilience as "managing risk by crisis."

What is still missing in many systems is a planning core that is physically stronger, more coherent across scales, and more directly translatable into decision-grade action. That is precisely what τ — understood as a law-faithful agro-climate-resilience twin — would provide.

Seven numbered key findings:

1. **The forecasting bottleneck is real and consequential.** Existing anticipatory-action programmes already work; their documented limitation is forecast quality, trigger calibration, and lead-time confidence rather than institutional design. Better seasonal physics directly addresses the binding constraint.
2. **A bounded-error planning twin changes the nature of the product.** The shift is not from a worse probabilistic outlook to a better one. It is from probabilistic scenario bundles to a structurally constrained action envelope — one that authorizes narrower, earlier, more confident decisions about planting, procurement, reserves, and social-protection triggers.
3. **The ROI baseline is already very high.** Nepal's WFP anticipatory action returns USD 34 per dollar compared with late response. Bangladesh anticipatory support cut the proportion of people going a day without eating by 36 percent. τ's contribution is improving trigger hit rates and targeting accuracy, compounding on this already-strong foundation.
4. **Six distinct opportunity clusters are addressable without new institutional architecture.** Seasonal production planning, anticipatory action triggers, food-system logistics and reserves, shock-responsive social protection, food-security dashboards, and last-mile climate-services delivery can all absorb τ outputs through existing pipelines — WMO RCOFs, WFP anticipatory programmes, World Bank FS-CAPs, and national advisory systems.
5. **The competitive landscape confirms the gap.** FEWS NET, GIEWS, WFP VAM/ALIVE, IRI Columbia, and EU MARS/JRC together represent mature institutional capacity in food-security early warning and seasonal outlooks. None provides a physics-grade, bounded-error planning twin that connects seasonal forecast uncertainty directly to decision-grade action triggers at district resolution.
6. **The finance architecture already exists.** CERF anticipatory action, GCF climate resilience windows, START Network, ECHO DIPECHO, and World Bank PROGREEN/Food Systems 2030 all provide eligible financing channels. Country-level τ integration costs are estimated at USD 4–10 million per country protecting 1–3 million people; regional corridor platforms at USD 25–60 million serving 30–50 million people.
7. **This is a protection story, not only a forecast story.** The highest-value consequences of stronger seasonal physics are not marginally better outlooks. They are avoided distress sales, reduced herd depletion, better-timed seed and fodder distribution, earlier reserve pre-positioning,

more credible social-protection triggers, and — ultimately — fewer preventable escalations from climate shock to food crisis.

2 Why This Matters Now

2.1 Compound shocks are rewriting the food-security risk landscape

The global food-security crisis of 2024 was not a single-cause event. Conflict, displacement, economic disruption, and climate extremes interacted in ways that overwhelmed systems built to respond to linear, single-hazard sequences. The Global Report on Food Crises 2025 documents more than 295 million people in acute hunger across 53 countries and territories — a figure that reflects not just production shortfalls, but the cascading failure of food-access systems under compound stress.

Climate-related hazards are now the primary driver of agricultural loss at the global level. FAO's 2025 disaster assessment documents USD 3.26 trillion in total agricultural losses from disasters between 1991 and 2023, with nearly USD 2.9 trillion — approximately 89 percent — attributable to climate hazards. The annual average of USD 99 billion in losses is a lower-bound estimate; it does not fully capture secondary losses through markets, nutrition, labor, and long-run investment suppression.

The structure of climate-related agricultural losses has also changed. The hazard profile is no longer predominantly acute — a single flood, a single drought — but increasingly compound and sequential: a delayed onset followed by a mid-season dry spell, then a closing-period flood, each interacting with already-stressed input and market systems. This compound profile defeats response architectures built around single-trigger, single-hazard logic.

2.2 Anticipatory action is scaling — and its binding constraint is forecast quality

The global anticipatory-action community has achieved something remarkable in the past decade. WFP now operates anticipatory-action programmes in 44 countries. Over 6 million people were protected in 2024 alone. The documented performance record — USD 7 in benefits per dollar invested, USD 34 in Nepal, 36 percent fewer people going without food in Bangladesh, 28 percent reduction in harmful coping strategies in Ethiopia — establishes anticipatory action as one of the highest-return humanitarian investments available.

But the system's practitioners are candid about its constraint. Forecast-based triggers depend on forecast quality and trigger calibration. False triggers waste scarce financing and erode institutional trust. Missed triggers leave populations exposed during the precise windows when pre-crisis action was most available. Lead-time confidence — the ability to act weeks or months before a shock peaks — is limited by the fidelity of the underlying sub-seasonal-to-seasonal outlook. The operational question is not whether to do anticipatory action; it is how to do it more reliably.

2.3 The food crisis 2024 baseline defines the scale of the problem

The 2024 acute-hunger figure of 295 million people is a planning datum, not a static endpoint. Behind it lie well-documented structural patterns: agricultural systems where planting and harvest timing is determined by late, poorly calibrated seasonal advisories; humanitarian response systems that are chronically too slow to pre-position because the relevant forecasts arrive too late or with too little confidence; and social-protection systems that scale up after crisis is visible rather than before it materializes.

The World Bank's food-security analysis frames this plainly: too many systems are still "managing risk by crisis." The available tools — early warning, rural infrastructure, climate-smart agriculture, safety nets, emergency planning — are individually present in most countries; what they lack is integration under a shared, physically coherent planning core that can authorize earlier, better-targeted action.

2.4 The institutional moment is aligned

WMO's agricultural-meteorology programme explicitly distinguishes the planning horizon (where climate information is most used) from the operational horizon (where recent weather and short-range forecasts are most used). The sub-seasonal-to-seasonal gap — roughly 10 to 90 days — is where planning decisions about seed procurement, planting windows, reserve policy, and anticipatory action are made, and where forecast skill historically degrades fastest. WMO's current sub-seasonal-to-seasonal agriculture project is specifically designed to close this gap for water-resource managers and food-security actors.

The τ opportunity is thus not a speculative institutional invention. It is an attempt to strengthen an architecture that already exists, with documented public-good demand and active institutional investment. The question is whether the underlying physics can be made sufficiently faithful to support decision-grade planning products at the relevant horizons.

3 Scope and Reader Orientation

This document is **Paper 4 of 5** in the Panta Rhei Agriculture Impact Portfolio. It focuses on the planning-and-resilience layer of the food system: the interface between seasonal and sub-seasonal climate science, anticipatory action, food-system logistics and reserves, and national and regional resilience planning.

In scope for Paper 4:

- Sub-seasonal and seasonal agricultural planning (crop calendars, planting windows, variety selection, fodder planning);
- Disaster anticipation and early action (forecast-based triggers, anticipatory cash, seed, and input distributions; flood and drought activation protocols);
- Food-system resilience at district, national, and regional levels (reserve policy, procurement pre-positioning, market continuity, social-protection triggers);
- Climate-service delivery into agriculture, social protection, and emergency planning;
- The governance and institutional interface between climate producers and food-system decision-makers.

Out of scope for Paper 4 (covered elsewhere in the portfolio):

- Paper 1: short-range operational agro-weather intelligence (day-to-day field operations, harvest timing, frost and heat warnings);
- Paper 2: irrigation management, soil moisture monitoring, and water-productivity optimization;
- Paper 3: pest, disease, and livestock-stress early warning;
- Paper 5: crop biology, photosynthesis engineering, breeding, and targeted gene design.

Primary audiences: Ministries of agriculture and planning; national meteorological and hydrological services; food-security agencies and IPC coordination bodies; disaster-risk reduction authorities; social-protection ministries; regional climate centres; WFP, FAO, and OCHA country operations; public development banks; strategic grain-reserve authorities; agricultural insurers; farmer organizations and cooperatives; and climate-agriculture funders.

4 The Opportunity Baseline

4.1 USD 99 billion per year in agricultural disaster losses

FAO's 2025 global disaster assessment provides the anchor figure: disasters caused USD 3.26 trillion in agricultural losses over the 1991–2023 period, averaging approximately USD 99 billion per year. Climate-related hazards — droughts, floods, storms, extreme temperatures, and their compound sequences — account for USD 2.9 trillion of that total, or roughly 89 percent. These losses fall disproportionately on low- and middle-income countries, and disproportionately on smallholder farmers, pastoralists, and subsistence producers with the least capacity to recover.

This baseline has direct planning implications. A significant portion of these losses is not meteorologically inevitable; it is planning-sensitive. Losses from delayed planting, mistimed reserve releases, under-prepared storage, late veterinary or fodder intervention, and reactive rather than anticipatory social protection are all, in principle, reducible through better seasonal planning — if the planning inputs are physically faithful enough to authorize earlier, more targeted decisions.

4.2 295 million people in acute hunger in 2024

The Global Report on Food Crises 2025 documents the human scale of the problem. More than 295 million people across 53 countries and territories experienced acute hunger — IPC Phase 3 or above — in 2024. Climate extremes were among the three major drivers, alongside conflict and economic disruption. The GRFC documents a multi-year pattern in which food crises in drought-affected regions of the Sahel, the Horn of Africa, and Southern Africa follow predictable climate sequences that, in principle, are addressable through better seasonal anticipation.

The 295 million figure is not a fixed ceiling. It reflects the current state of food-system protection under current planning tools. Marginal improvements in trigger quality, lead-time confidence, and pre-positioning accuracy — the precise capabilities a τ-grade planning twin would improve — have potential impact at scale even without changing the broader political or economic drivers of food insecurity.

4.3 WFP's anticipatory-action benchmark: USD 1 in, USD 7 out

WFP's anticipatory-action documentation provides perhaps the clearest ROI benchmark in humanitarian practice. Across its portfolio:

- Every USD 1 invested in anticipatory action yields up to USD 7 in avoided losses and added benefits;
- In Nepal, WFP reports USD 34 saved per dollar compared with late response;
- In Bangladesh, anticipatory support triggered ahead of floods reduced the proportion of people going a day without eating by 36 percent;
- In Ethiopia, families receiving early action combined with cash transfers reduced harmful coping strategies by 28 percent, outperforming cash-only interventions.

These figures establish that the institutional architecture already works. The strategic question is not whether anticipatory action is worth investing in; it is how to extend its reach, sharpen its triggers, and reduce the false-trigger and missed-trigger rates that currently constrain its scaling.

WFP currently covers 44 countries with anticipatory-action programming. The programme scope is explicitly limited by financing, but operationally it is also constrained by forecast quality. Improving the physics of the underlying seasonal outlook is thus a direct force multiplier on an already-proven investment.

4.4 WFP's 44-country portfolio and the trigger-quality constraint

WFP's current anticipatory-action programmes link pre-agreed financing to pre-agreed forecast thresholds. When a climate indicator — flood-forecast exceedance, drought probability, cyclone track — crosses a trigger threshold, pre-positioned resources are released. The architecture is sound; its documented limitation is that forecast thresholds are set under uncertainty about the true underlying seasonal probability, which means triggers are calibrated conservatively to manage false-activation risk. This conservatism is prudent but costly: it means some genuine threats are missed or acted upon too late.

4.5 The \$1 → \$7 ROI and the forecast-quality multiplier

The USD 1 → USD 7 ROI benchmark implicitly assumes current trigger quality. Under τ, the realistic-optimistic claim is not that these benefits suddenly appear from nowhere — the institutional infrastructure already produces them. The claim is that improved trigger hit rates, better district-level targeting, and extended actionable lead time compound the existing benefit-to-cost ratio. If trigger accuracy improves enough to reduce false triggers by 30 percent and missed triggers by 20 percent, the effective benefit-to-cost ratio could approach or exceed the Nepal benchmark of USD 34 for programs currently operating at USD 7.

5 Working τ Assumptions

This paper adopts an explicitly conditional stance. It asks what practical and public-good consequences would follow if the τ framework's relevant claims were sufficiently accurate for planning use. The following assumptions are made for analytical purposes and do not constitute claims that the broader scientific or agricultural communities have accepted them.

5.1 Climate and hazard assumptions

- τ provides a bounded-error, constructive, coarse-grainable representation of sub-seasonal to seasonal atmospheric, hydrological, and land-surface dynamics.
- τ outputs are physically faithful in their handling of rainfall onset and cessation, dry-spell risk, seasonal flood envelopes, temperature anomaly sequences, soil-moisture evolution, and multi-stage compound hazard structures.
- Uncertainty bounds in τ outputs are structurally derived — reflecting the genuine information content of the physical system at the relevant horizon — rather than ensemble spread calibrated post hoc.
- τ handles both rapid-onset hazards (flash floods, cyclones) and slow-onset hazards (drought development, seasonal rainfall failure, crop-water stress accumulation) within a unified framework.

5.2 Agrifood-system translation assumptions

- τ outputs can be expressed as district- and national-level planning products — action envelopes — rather than only field-level point advisories.
- τ supports planning across crop calendars, pasture and fodder cycles, seed and input delivery timing, storage and reserve strategy, social-protection trigger design, logistics and procurement, and anticipatory humanitarian activation.
- τ planning outputs are designed to plug into, rather than replace, existing climate-service and early-warning architectures: RCOFs, national climate forums, IPC coordination processes, and food-security dashboards.

5.3 Food-security and policy translation assumptions

- Governments, humanitarian actors, and food-system planners can use τ products to move from broad climate warning to decision-grade action — with enough confidence to justify pre-crisis resource deployment.
- τ can support the specific planning exercises currently pursued through WMO regional and national climate forums, World Bank FS-CAPs, FAO early warning / early action frameworks, WFP anticipatory-action programmes, and strategic reserve planning.

5.4 What this paper does not assume

This paper does not assume that ministries will immediately adopt the τ ontology wholesale, that political and institutional bottlenecks will dissolve with better forecasts, or that climate-science improvement alone resolves the conflict, displacement, and humanitarian-funding crises that are co-drivers of food insecurity. The claim is practical and bounded: if the planning twin is physically stronger, a substantial and identifiable class of avoidable losses and delayed actions becomes tractable.

6 What Changes with a Law-Faithful Twin

6.1 From probabilistic outlook to bounded action envelope

Today's seasonal planning architecture typically has the following structure: a climate outlook is produced, usually as a probabilistic tercile map or seasonal anomaly projection; sector experts interpret it for agriculture, water, or humanitarian use; advisories and contingency notes are produced; and implementation then proceeds under large residual uncertainty about timing, geographic concentration, intensity, and duration.

Under the strongest τ assumption, this chain is fundamentally restructured. The planning product becomes not a probabilistic statement that conditions in the next season are likely to be above, near, or below normal — it becomes a structurally bounded action envelope: a specification of the range of seasonal trajectories consistent with the physical laws governing the system at that forecast horizon, at that spatial resolution, given the current initial state.

The difference is not semantic. A probabilistic tercile map tells a ministry of agriculture that there is a 60 percent chance of below-normal rainfall. A bounded action envelope tells it that, at district resolution, the rainfall trajectory will stay within a calculable range with a calculable degree of confidence, and that the lower tail of that range implies a specific pattern of crop-water stress emergence beginning in a specific week — with an error bound that is structurally honest rather than heuristically calibrated. The second product licenses earlier, narrower, more consequential decisions.

6.2 Trigger quality and anticipatory action

The anticipatory-action community's core operational challenge is trigger calibration: setting thresholds that are sensitive enough to catch genuine threats without generating false activations that drain trust and resources. Current triggers are calibrated under uncertainty about the seasonal probability distribution, which forces conservative threshold-setting.

A τ-grade planning twin changes the calibration problem. If uncertainty bounds are structurally derived and physically honest, trigger design can be grounded in the actual information content of the forecast at the relevant lead time rather than in ensemble-spread estimates calibrated from historical verification. This supports:

- Earlier triggering with higher confidence at longer lead times;
- Narrower geographic targeting, reducing wasted resource deployment in unaffected areas;
- Better package design — distinguishing flood-preparedness packages from drought-preparedness packages at the district level rather than deploying generic packages across wide regions;
- More defensible trigger logic for pre-agreed financing mechanisms, reducing the governance risk of post-hoc challenge.

6.3 Food-system resilience as an upstream function

The dominant mode of food-system resilience today is reactive: reserves are released after prices spike; social-protection transfers scale up after crisis is confirmed; humanitarian pre-positioning begins after early-warning alert levels are elevated. These responses often arrive too late to prevent the most costly outcomes — herd depletion, distress sales, child malnutrition peaks, seed stock exhaustion.

A τ-grade planning twin makes resilience an upstream function rather than a downstream repair. Stock and reserve sizing can be calibrated against the seasonal envelope rather than against last year's crisis experience. Procurement and pre-positioning can begin when the bounded probability of a bad season crosses an action threshold, rather than when the bad season is confirmed. Social-protection registers can be updated and transfer-readiness achieved before the shock peaks, not after. Fodder and veterinary inputs can be positioned in pastoral corridors based on modeled seasonal water-point and pasture trajectories.

6.4 Cross-institutional bridge function

One of the most important but least visible consequences of a physically coherent planning twin is institutional: it provides a shared physical basis on which organizations that currently operate separately can coordinate their responses. Meteorological agencies, ministries of agriculture, water authorities, grain-reserve bodies, social-protection ministries, and humanitarian actors all use different information products built on different physical assumptions and calibrated against different historical datasets.

A τ-grade planning twin — if adopted as a shared reference — provides a common physical object against which all these actors can align their threshold decisions. That coordination function alone, independent of improved forecast skill, has large potential value in systems where the current resilience gap is not a lack of data but a lack of joined-up action.

6.5 Coarse-grained national and regional twins

If the τ substrate genuinely supports certified coarse-graining — the ability to aggregate district-level outputs to national and regional scales while preserving the structural integrity of uncertainty bounds — then the national food-system planning twin becomes a decision instrument rather than a visualization dashboard. That implies:

- Reserve policy can be tied to ensemble-informed seasonal scenarios with quantified confidence;
- Regional procurement planning can begin when regional seasonal outlooks cross pre-agreed thresholds;
- Early-warning systems can transition from monitoring to pre-activation mode at structurally justified moments rather than on the basis of committee judgment.

This is potentially one of the most transformative medium-term consequences of the τ agriculture stack.

7 Competitive and Incumbent Landscape

Five major institutional programmes currently anchor the global food-security early warning and seasonal planning architecture. Each represents genuine capability and long-run institutional investment. Each also has a specific and well-understood limitation that τ, under the stated assumption, would address.

7.1 FEWS NET — Famine Early Warning Systems Network

What it does well. FEWS NET is USAID's flagship food-security early warning system, operating in more than 30 countries across sub-Saharan Africa, Central Asia, and Latin America and the Caribbean. It produces integrated food-security condition and outlook reports using the IPC framework, with strong capacity for scenario-based analysis, livelihood analysis, and policy framing. FEWS NET's IPC Phase 3+ alerts have triggered institutional response in multiple countries. Its cadence of monthly updates and quarterly scenario reports is the standard reference for humanitarian and government planning in the countries it covers.

What it lacks. FEWS NET is a food-security analytical system, not a physics-grade forecast twin. Its seasonal outlooks rely on inputs from NMHSs, IRI, NOAA, ECMWF, and other providers — they synthesize existing products, they do not generate physics-grade bounded-error seasonal probability envelopes independently. The temporal cadence (monthly updates, quarterly outlooks) is slower than anticipatory-action trigger protocols require. District-level precipitation probability and crop-water stress precision are limited by the quality of the underlying seasonal forecast inputs.

How τ would differ. τ would operate at a sub-seasonal physics layer below FEWS NET's analytical synthesis layer, providing higher-resolution, bounded-error seasonal forecast inputs that could improve the IPC scenario calibration and allow FEWS NET's analytical capacity to be applied to sharper, more actionable physical data.

7.2 GIEWS — FAO Global Information and Early Warning System

What it does well. GIEWS is the UN's official global food-security and crop-situation monitoring system, with country-level coverage across more than 100 nations. It produces crop monitoring reports, food-security country briefs, and supply-and-demand assessments that form the statistical backbone of FAO's global food-security analysis. It is updated two to three times per year per country at the national level, with faster-turnaround emergency alerts for acute crises. GIEWS is authoritative, globally consistent, and institutionally embedded in UN processes.

What it lacks. The two-to-three times per year update cadence, while appropriate for aggregate supply monitoring, is too slow for operational anticipatory-action trigger decisions. GIEWS' food-security coverage is strong at national and regional scales; sub-national and district-level operational precision is limited. Like FEWS NET, GIEWS synthesizes available seasonal forecast inputs rather than generating bounded-error planning-grade forecast products independently.

How τ would differ. τ would operate at a faster cadence and higher spatial resolution in the seasonal forecast domain, providing inputs that could improve the timeliness and precision of GIEWS crop-situation assessments and close the gap between national-level monitoring and district-level anticipatory action.

7.3 WFP VAM / ALIVE — Vulnerability Analysis and Mapping and ALIVE Platform

What it does well. WFP's VAM unit and its ALIVE platform are the operational backbone of WFP's food-security targeting and trigger decisions. VAM integrates food-security, market, livelihoods, and shock data to produce the vulnerability maps and analytical products that inform WFP programming decisions. ALIVE is designed to make VAM's outputs more interoperable and machine-readable for trigger-based financing and anticipatory action. It is the closest existing system to an operational anticipatory-action decision engine.

What it lacks. VAM's strength is in food-security and livelihood integration — it aggregates inputs from weather, markets, and field surveys into operational vulnerability products. Its forward seasonal physics layer relies on the same probabilistic climate products as the broader system. It does not generate structurally bounded seasonal forecast envelopes; it consumes them. The quality of trigger calibration is therefore constrained by the quality of the forecast inputs.

How τ would differ. τ would provide higher-confidence forecast inputs into ALIVE's trigger logic, improving the physical grounding of activation thresholds without requiring changes to ALIVE's broader integration and targeting architecture.

7.4 IRI Columbia — International Research Institute for Climate and Society

What it does well. IRI is the global academic leader in seasonal probabilistic climate forecasting and climate-information services for decision-making. Its ENSO forecasts, seasonal outlook products, and capacity-building work with NMHSs are foundational to the global climate-services architecture. IRI's seasonal forecasting skill at the 1–3 month horizon is among the best available globally. Its work on translating seasonal climate information for food-security, agricultural, and health applications is technically rigorous and institutionally respected.

What it lacks. IRI's products are probabilistic seasonal outlook products — tercile probability maps and ensemble-based anomaly forecasts. They do not provide physics-grade bounded-error planning envelopes with structurally derived uncertainty bounds at district resolution. The translation from IRI's seasonal probabilistic output to operational action trigger — what to do when the 3-month forecast shows a 60 percent chance of below-normal rainfall — remains a significant source of trigger-calibration uncertainty in operational anticipatory-action programmes.

How τ would differ. τ's most direct point of engagement with IRI's architecture is at the physics layer: providing a complementary bounded-error approach alongside IRI's ensemble-probabilistic approach, potentially enabling a physical consensus between the two methods for operational trigger calibration.

7.5 EU JRC / Copernicus MARS — Agricultural Monitoring Service

What it does well. The JRC Monitoring Agricultural ResourceS (MARS) unit, operating through the EU's Copernicus Land Monitoring Service, provides Europe-wide and global satellite-based crop-monitoring and yield-forecasting products. MARS produces monthly crop-monitoring bulletins for Europe and, through GEOGLAM/AMIS, contributes to global agricultural market monitoring. Its satellite-derived indicators — NDVI, evapotranspiration, soil-moisture anomalies, crop-yield estimates — are technically sophisticated and feed into EU and global food-security assessments.

What it lacks. MARS is a retrospective and near-real-time crop-monitoring system, not a forward seasonal planning twin. It characterizes crop conditions and yield trajectories as they unfold; it does not generate bounded-error forward seasonal risk envelopes for planning and anticipatory action. Its global coverage is strongest for large-scale commercial crop systems; coverage of smallholder-dominated sub-Saharan Africa and South Asian systems is less operationally precise. Shock anticipation and action-trigger design are outside its core mandate.

How τ would differ. τ would operate in the forward planning dimension that MARS does not cover: generating the seasonal drought and crop-water stress probability envelopes that anticipatory-action systems require, which MARS outputs could then validate retrospectively as seasons unfold.

8 Structured Opportunity Map

8.1 Cluster A — Seasonal Production Planning

8.1.1 A1. Crop and variety selection under seasonal outlook envelopes

The most consequential planning decision in rainfed agriculture is made before the season begins: what to plant, at what scale, in which locations, and with what risk tolerance. Under current planning tools, this decision is made against probabilistic tercile outlooks that cannot confidently discriminate between a drought-shortened season and a near-normal season at district resolution.

Under τ, seasonal variety and crop-type decisions could be grounded in bounded-error production-risk envelopes at district resolution: a physically derived range of season-length, rainfall-amount, and crop-water-stress scenarios, each with calculable probability bounds. This would support more precise guidance on drought-tolerant versus standard variety allocation, on early-maturing versus full-season crop selection, and on area allocation between rainfed and irrigated production zones.

8.1.2 A2. Planting-window and campaign sequencing

Optimal planting windows are sensitive to rainy-season onset timing, which is one of the most planning-relevant but forecast-uncertain quantities in tropical agriculture. Early planting in a false-start season leads to seedling mortality and replanting costs; late planting after a delayed onset shortens the effective growing season and reduces yield potential. Both error types are costly, and both are addressable with better bounded-error onset probability.

Under τ, district-level onset probability envelopes could support crop-campaign sequencing — staggered planting across zones based on the bounded onset distribution, second-crop decisions contingent on first-season performance forecasts, and early-season input delivery phased against the onset envelope. This is distinct from the operational planting-window advisory of Paper 1; Paper 4 addresses the campaign-level investment and logistics decisions made weeks before planting begins.

8.1.3 A3. Pasture, fodder, and mixed-system seasonal planning

Pastoral and agropastoral systems are disproportionately exposed to seasonal planning failure. Herd management decisions — when to move, where to graze, when to sell, when to destock — are made against long-range expectations about pasture and water-point availability that current probabilistic outlooks cannot support at the resolution and lead time needed for operational decisions.

WMO's ACREI programme in the Greater Horn of Africa provides the operational template: downscaled forecasts and advisories to help farmers and pastoralists decide when, what, where, and how to plant and manage land, crops, pastures, water, and preserve food and feed as the season progresses. Under τ, the physical inputs to this advisory process could be substantially more precise, supporting earlier decisions about seasonal migration routes, fodder-reserve sizing, and water-point preparation.

8.2 Cluster B — Anticipatory Action and Disaster-Risk Reduction

8.2.1 B1. Forecast-based drought and flood activation

WFP's anticipatory-action architecture provides the operational template: pre-agreed financing linked to pre-agreed forecast thresholds, releasing pre-positioned resources before shocks peak. The current constraint is trigger calibration under seasonal forecast uncertainty.

Under τ, trigger design can be grounded in structurally bounded probability envelopes: the activation threshold is set at the point where the bounded seasonal trajectory enters the range consistent with historical damage profiles, rather than at an empirically calibrated ensemble-probability threshold. This change — subtle in institutional terms, large in operational terms — has the potential to:

- Increase early-activation rates without increasing false-activation rates, by distinguishing more reliably between genuine seasonal deterioration and normal variability;
- Extend actionable lead time from the current 10–30 day effective window for many programmes to 30–90 days, covering the period where pre-positioning and procurement decisions can still be made;
- Reduce the area of activation packages, by resolving the spatial structure of risk more precisely at district scale.

8.2.2 B2. Anticipatory agricultural support packages

Before-shock agricultural support — seed, feed, veterinary inputs, cash for fodder purchase, water-trucking pre-positioning, emergency irrigation support — has a narrow effectiveness window. Once herds are depleted, seed stocks are consumed, or planting windows are missed, the damage cannot be undone in-season. The value of earlier, more targeted anticipatory agricultural packages is thus asymmetric: acting two weeks early is vastly more cost-effective than acting two weeks late.

Under τ, package design can move from generic early-action bundles to differentiated, district-specific packages calibrated to the bounded trajectory of the modeled stress pathway: flood-preparedness packages for river-basin zones entering a high-probability flood envelope, drought packages for districts entering a low-rainfall trajectory, combined seed-and-fodder packages for agropastoral zones showing both rainfall-deficit and pasture-stress signals.

8.2.3 B3. National contingency planning tied to forecast thresholds

WFP's work in Mozambique demonstrates that anticipatory actions can be embedded in national contingency plans and tied to drought triggers linked to social-protection scale-up. Under τ, national contingency plans can be designed around physically bounded seasonal probability thresholds rather than around historical recurrence intervals calibrated from short observational records. This matters particularly in countries where climate trends are already shifting historical frequencies, making past-recurrence-based trigger design systematically miscalibrated.

8.3 Cluster C — Food-System Logistics, Storage, and Reserve Intelligence

8.3.1 C1. Seasonal storage and market-continuity planning

Seasonal shocks disturb not only production but the entire food-system logistics chain: storage utilization, transport availability, market arrival timing, and price behavior all shift in ways that compound production-side damage. Pre-positioning storage capacity, pre-clearing transport corridors, and pre-contracting market-continuity arrangements all have value that is inversely related to lead time — shorter lead time means higher cost and lower coverage.

Under τ, the seasonal logistics planning layer can be activated earlier and at higher confidence, allowing national authorities and humanitarian logistics systems to begin storage preparation and market-continuity planning before the shock is confirmed rather than as it arrives.

8.3.2 C2. Strategic grain reserve timing and release planning

The 2025 World Bank/WFP/FAO strategic grain reserve review provides important calibration for this cluster. It argues that strategic grain reserves should remain relatively small, with procurement timing optimized and release policy designed to minimize fiscal cost and market distortion. The implication is that reserve policy decisions are extremely sensitive to the seasonal risk outlook: whether to replenish, when to replenish, when to hold, and when to release, all depend on the accuracy of the forward seasonal food-availability and import-cost picture.

Under τ, reserve policy can become scenario-informed rather than reactive:

- Replenishment decisions can be timed to the predicted seasonal production trajectory, buying on import markets before domestic shortage drives prices up;
- Release decisions can be calibrated to the bounded probability of a market-failure event, avoiding premature releases that are difficult to reverse;
- Regional coordination on reserve sharing and mutual drawdown can be authorized against common bounded seasonal outlooks rather than conflicting national estimates.

8.3.3 C3. Procurement and pre-positioning for humanitarian response

WFP, FAO, and national humanitarian systems make procurement and logistics pre-positioning decisions whose lead times (often 6–12 weeks for international procurement, 4–8 weeks for domestic procurement and in-country movement) frequently exceed the current practical actionable lead time of seasonal outlook products. The result is chronic under-positioning: humanitarian stocks are frequently still in procurement when shocks peak.

Under τ, procurement authorization decisions could be moved to the period when the bounded probability of a shock event crosses the pre-agreed action threshold — potentially 30–60 days earlier than current practice — substantially improving the match between in-country stock levels and realized demand at shock onset.

8.4 Cluster D — Social Protection, Insurance, and Risk Finance

8.4.1 D1. Shock-responsive social protection triggers

Shock-responsive social protection — the capacity to expand cash-transfer and in-kind support programmes rapidly in response to climate shocks — has become a core component of food-system resilience architecture in many countries. The operational constraint is trigger timing: activating scale-up before crisis is visible means acting on forecast signals rather than confirmed damage, which requires forecast confidence that current seasonal outlook products often cannot provide.

Under τ, social-protection trigger logic can be grounded in bounded-error seasonal probability envelopes, authorizing automatic scale-up preparation when the bounded trajectory enters a high-risk range, rather than waiting for confirmed crisis activation. The WFP work in Mozambique and the Dominican Republic provides a direct operational template; τ improves the physical inputs without requiring institutional redesign.

8.4.2 D2. Insurance and disaster-risk finance aligned to forecast quality

Parametric agricultural insurance — where payouts are triggered by meteorological indices rather than damage assessment — has significant scaling potential but is constrained by basis risk: the divergence between the index trigger and actual farm-level losses. Basis risk is directly related to the precision of the underlying climate-index products. More physically faithful seasonal climate products reduce basis risk, improving insurance actuarial performance and making insurance more accessible to smallholders.

WFP-supported disaster-risk finance solutions have reached 6.2 million people since 2017. Under τ, the index quality underlying parametric trigger products could be improved enough to reduce basis risk materially, supporting wider smallholder uptake without increases in premium costs.

8.5 Cluster E — Food-Security Dashboards, National Planning, and Public Investment

8.5.1 E1. FS-CAPs and climate-smart investment prioritization

The World Bank's Food System Climate Action Plans provide a strategic vehicle for integrating climate-risk analysis into national food-system investment and policy planning. Under τ, the physical evidence base beneath FS-CAPs could be strengthened: bounded-error seasonal risk projections could support more precisely targeted infrastructure investments, more defensible safety-net readiness standards, and more risk-grounded crop-system transition strategies.

8.5.2 E2. National food-security dashboards and early-warning integration

The planning need for national food-security dashboards is not weather data alone; it is connected food-system intelligence — seasonal outlook, water availability, crop performance, market stress indicators, reserve position, and safety-net readiness, all updated at consistent frequency and integrated on a shared physical basis. Under τ, the seasonal outlook layer of such dashboards could become physics-grade rather than expert-curated, improving the timeliness, granularity, and confidence of the overall planning picture.

8.5.3 E3. Risk-informed infrastructure prioritization

World Bank food-security analysis highlights rural infrastructure — roads, storage facilities, irrigation infrastructure, and local market systems — as a structural resilience priority. Under τ, infrastructure investment prioritization could be grounded in bounded-error seasonal risk maps: which districts face the highest probability of recurrent drought, which river-basin zones face the highest multi-year flood frequency, and where storage and logistics buffers would yield the highest marginal reduction in seasonal food-system fragility.

8.6 Cluster F — Climate Services and Last-Mile Planning Delivery

8.6.1 F1. PICSA-style extension-mediated planning support

WFP's climate-services work in Haiti demonstrates what last-mile delivery of seasonal planning information can achieve: more than 6,500 farmers received PICSA-based seasonal planning support using weather data and seasonal forecasts to guide planting and harvesting decisions. In Zimbabwe, WFP supports locally tailored climate advisories in local languages in ten-day formats.

Paper 4 extends this logic from the farm/household level toward the district and sub-national food-system planning layer. The institutional template already exists; τ would improve the quality of the seasonal inputs flowing into it.

8.6.2 F2. National and regional climate-forum integration

WMO's Regional Climate Outlook Forums and national follow-on forums are the established institutional bridge between climate producers and sector users. τ could strengthen the technical substance inside that bridge — providing more precise, bounded-error seasonal inputs that RCOFs can translate into sharper, more actionable advisories for agricultural extension, disaster-risk reduction, and food-security planning. This does not require replacing the RCOF institutional architecture; it requires improving the physical quality of what flows through it.

9 Geographic Case Studies

9.1 Case Study 1: WFP Anticipatory Action — Bangladesh and Nepal Pilots

Context and baseline performance. WFP's anticipatory-action programmes in Bangladesh and Nepal are among the most documented and analytically mature in its global portfolio, and they provide the clearest evidence of both the potential and the constraints of forecast-based humanitarian action.

In Bangladesh, WFP operates anticipatory-action programming covering approximately 500,000 to 1 million people in flood-exposed districts, linked to trigger thresholds from the Bangladesh Flood Forecasting and Warning Centre (FFWC). When assistance was activated ahead of floods, WFP documented that 36 percent fewer people went a day without eating during the flood event, compared with matched communities without anticipatory support. This is a substantial protection outcome from a single pre-activation cycle.

In Nepal, WFP's anticipatory-action programme has supported multi-hazard pre-activation, including cash transfers before monsoon flooding in flood-prone districts. An independent return-on-investment analysis documented that every dollar invested before a disaster in Nepal saved up to USD 34 compared with responding after the event. That figure reflects not only the direct value of cash and in-kind support, but the avoided losses from distress sales, herd depletion, and emergency health expenditures that are prevented when transfers arrive before the shock peak.

Forecast-based trigger constraints. Both programmes illustrate the operational constraint clearly. Flood-forecast-based triggers in Bangladesh rely on FFWC ensemble forecasts and gauge-based river-level thresholds. These triggers are accurate at 3–7 day lead times for the immediate flood-arrival signal, but the effective planning window for pre-positioning cash, arranging beneficiary communication, and organizing distribution logistics is considerably longer than 7 days. At 10–30 days lead time — the operational planning window — flood-probability forecasts lose enough resolution that trigger calibration must be conservative to avoid false activations.

The consequence is that some activations arrive late relative to the shock — after flood peaks are already developing — and some potential activations are not triggered because the seasonal probability signal was not strong enough at the required lead time to authorize pre-positioned spending.

τ-enabled change. Under the τ assumption, the Bangladesh and Nepal programmes would benefit from:

- More physically faithful 30–90-day flood and drought probability envelopes at district resolution, extending the actionable planning window from roughly 7 days to 30–60 days for major flood events;
- Better trigger calibration grounded in structurally bounded probability rather than ensemble-spread estimates, reducing false-trigger rates without increasing missed-trigger rates;
- Integration of seasonal food-access modeling — connecting the flood-risk envelope to household

food-access trajectories based on modeled crop and market effects — allowing more precise package design (which households need cash-plus-seed versus cash-only, for instance);
- Earlier procurement and distribution-logistics authorization, reducing the emergency-premium cost of last-minute procurement and improving the match between beneficiary lists and population-at-risk.

The WFP Bangladesh model already interfaces with government social-protection systems and the FFWC institutional structure. τ integration would operate at the forecast-input layer, providing improved physical inputs to the existing trigger architecture rather than replacing it.

Reference organizations: WFP Bangladesh and Nepal country offices; Bangladesh Flood Forecasting and Warning Centre (FFWC); OCHA Bangladesh; World Bank social-protection team; Government of Bangladesh Ministry of Disaster Management and Relief.

9.2 Case Study 2: FEWS NET / WFP Southern Africa — El Niño 2023–24 Anticipatory Response

Context and El Niño onset. The 2023–24 El Niño event brought one of the most severe seasonal droughts to Southern Africa in the post-2000 record, with Zimbabwe, Zambia, Malawi, and Mozambique all experiencing significantly below-normal seasonal rainfall during the 2023–24 growing season. The Southern African Development Community (SADC) declared a regional humanitarian disaster in mid-2024, with approximately 20 million people across the region facing acute food insecurity.

Critically, the El Niño signal was scientifically detectable and publicly communicated more than six months before the drought fully materialized. NOAA, IRI, ECMWF, and WMO all issued El Niño forecasts and Southern Africa rainfall deficit warnings well in advance of the 2023–24 agricultural season. FEWS NET and GIEWS provided IPC Phase 3+ outlooks for Zimbabwe and Zambia by November 2023, before the harvest was complete.

Limits of the response. Despite advance warning, the humanitarian response in several countries was substantially reactive rather than anticipatory. WFP and FAO did pre-position some stocks and trigger limited anticipatory cash distributions in parts of the region, but the scale and timing of the response lagged behind the scale and timing of the need.

The primary constraint was precision: El Niño-based seasonal outlooks for Southern Africa are accurate at continental and sub-regional scale — most of Southern Africa will be drier than normal — but they lack district-level precipitation probability and crop-water stress precision sufficient to authorize major pre-crisis procurement and distribution decisions with confidence. Trigger calibration delays in several countries reflected the genuine difficulty of moving from “Southern Africa will be drier than normal” to “District X in Zambia will experience crop-failure-level rainfall deficit beginning in week 8 of the season with 70 percent probability.”

FEWS NET’s analytical capacity is high; the limitation is in the physical forecast layer beneath it. IPC Phase 3+ projections for Zimbabwe were accurate at national scale, but the district-level disaggregation of crop failure probability — needed to authorize district-specific anticipatory distributions — was not precise enough to close procurement, logistics, and social-protection decisions with confidence at the required lead time.

τ-enabled change. Under the τ assumption, the Southern Africa El Niño response would have been addressable with:

- District-level bounded-error seasonal drought probability envelopes for the major agroecological zones of Zimbabwe, Zambia, and Malawi, generated 8–12 weeks before the planting-window decision point (November–December 2023);
- Crop-water stress trajectory forecasting tied to planting-window advisories — distinguishing

districts where mid-season dry spells would be the primary risk from districts where season-end rainfall failure was the primary risk, with quantified probability bounds;
- Earlier procurement authorization for WFP and national government reserves: if the bounded seasonal trajectory for a district crosses the activation threshold in November, procurement authorization 10–12 weeks before peak stress allows domestic procurement at pre-crisis prices and avoids emergency-import premiums;
- Shock-responsive social-protection scale-up triggers set 6–10 weeks earlier than current practice, allowing social-protection registers to be updated, transfer mechanisms to be activated, and beneficiary communication to be completed before peak distress rather than during it.

The institutional infrastructure — FEWS NET's analytical system, SADC's early-warning coordination, WFP's Southern Africa anticipatory-action programmes, and FAO's GIEWS crop monitoring — already existed and was in use. What was constrained was the physical precision of the district-level seasonal forecast inputs, not the institutional architecture for translating warnings into action.

Reference organizations: FEWS NET; SADC Secretariat; WFP Southern Africa Regional Bureau; FAO GIEWS; OCHA Southern Africa; World Food Programme country offices in Zimbabwe, Zambia, Malawi, and Mozambique.

9.3 Case Study 3 (Supplementary): Greater Horn of Africa — ACREI / RCOF Integration

Context. The WMO/FAO Agricultural Climate Risk and Early Warning Integration (ACREI) programme in the Greater Horn of Africa provides one of the most explicit institutional models for translating seasonal climate information into agricultural advisory products at the level of household and community planning. ACREI was operational in 2023–2024, working within the WMO Greater Horn of Africa Climate Outlook Forum (GHACOF) process.

WMO's Regional Climate Outlook Forum for the Greater Horn of Africa provides a quarterly seasonal outlook; national forums then translate this into country-level agricultural sector advisories. The ACREI programme adds a layer of downscaling and crop-model integration, producing sub-national advisories on planting timing, variety selection, fodder planning, and water-point management for approximately 10 million pastoralists and smallholders in the region.

τ-enabled change. Under the τ assumption, ACREI's sub-seasonal input quality could be substantially improved. The key operational benefit is in the downscaling layer: replacing probabilistic tercile inputs — which arrive at GHACOF seasonal resolution — with bounded-error sub-seasonal inputs that resolve week-by-week and district-by-district probability trajectories. This would allow ACREI's crop-model and advisory layer to generate more precisely timed, more geographically differentiated advisories, improving the planting, fodder, and livestock-management guidance for smallholders and pastoralists in a system already designed to deliver it.

10 Finance, ROI, and Climate-Finance Eligibility

10.1 Existing finance architecture

The anticipatory-action and food-system resilience finance landscape is well-developed and already identifies forecast quality as an investable priority. Multiple streams are directly relevant to τ integration.

OCHA CERF Anticipatory Action Pillar. The UN Central Emergency Response Fund has committed approximately USD 35 million per year to its anticipatory action pillar, disbursed against pre-agreed trigger-based protocols before disasters strike. CERF anticipatory action is expanding rapidly; its 2023–2025 trajectory has seen roughly doubling of annual disbursement. τ integration into trigger design is directly eligible for CERF anticipatory-action funding as a component of forecast-based financing protocols.

WFP Anticipatory Action Funding. WFP's anticipatory-action programmes are financed through a combination of bilateral ECHO/DIPECHO funding, the START Network Anticipation Window, and WFP's own resilience resources. The START Network Anticipation Window specifically supports forecast-based financing for smaller, faster-response NGO-led anticipatory actions, with particular relevance for last-mile delivery in contexts where WFP country programmes are not the primary implementer.

ECHO DIPECHO. The European Commission's DIPECHO programme funds disaster preparedness and anticipatory action across WFP's targeted geographies. It is a major source of operational co-financing for forecast-based activation protocols and climate-service integration in food-crisis-prone countries.

Green Climate Fund (GCF) Climate Resilience Window. GCF provides funding for National Adaptation Plans and anticipatory action for food systems, with specific windows relevant to seasonal forecast improvement and early-warning integration. National Adaptation Plans in food-insecure countries frequently identify early-warning and seasonal forecast improvement as priority investable actions. τ integration into national climate services would be eligible under the GCF NAP support framework.

World Bank PROGREEN and Food Systems 2030. The World Bank's PROGREEN facility and its Food Systems 2030 initiative both support investments in climate-smart agriculture, food-system resilience, and early-warning system improvement. World Bank FS-CAP processes explicitly identify seasonal climate information quality as an investment priority, and the Bank's operational lending windows in agriculture and DRM are open to forecast-quality improvement as a bankable investment.

10.2 Cost scenarios

Scenario 1: Country-level τ integration into seasonal climate services and anticipatory-action trigger systems.

Scope: Integration of τ seasonal forecast outputs into one country's NMHS-linked seasonal outlook process, FEWS NET or GIEWS analytical pipeline, and WFP or government anticipatory-action trigger protocols. Includes technical capacity-building, trigger recalibration, and dashboard development.

Estimated cost: USD 4–10 million over three years.

Approximate protected population: 1–3 million people in the country's anticipatory-action coverage area.

This scenario is eligible for bilateral ECHO/DIPECHO, START Network, WFP anticipatory-action co-financing, and World Bank PROGREEN.

Scenario 2: Regional food-system resilience platform (Sahel or Southern Africa corridor).

Scope: Multi-country regional τ integration covering 10–12 countries, providing a shared bounded-error seasonal forecast layer for WFP's regional anticipatory-action programming, FEWS NET/GIEWS regional outlook integration, SADC/CILSS food-security coordination, and RCOF-linked advisory systems.

Estimated cost: USD 25–60 million over four years.

Approximate protected population: 30–50 million people in anticipatory-action coverage across the region.

This scenario is eligible for GCF climate resilience funding, World Bank Food Systems 2030, bilateral ECHO, and OCHA CERF mechanism alignment.

10.3 Benefit-to-cost analysis

WFP baseline ROI. WFP's documented USD 1 → USD 7 benefit-to-cost ratio across its global portfolio provides the reference baseline. Nepal's USD 1 → USD 34 represents the upper end of documented results under well-calibrated triggers. The gap between the average USD 7 and Nepal's USD 34 reflects in part the quality of trigger calibration and targeting precision in individual programmes.

τ contribution to ROI. τ's contribution to the benefit-to-cost ratio is not in replacing the anticipatory-action infrastructure but in improving its hit rate and lead-time performance. If trigger improvements reduce false-trigger rates by 25–35 percent and extend actionable lead time by 20–40 days:

- The cost per correctly triggered beneficiary falls (fewer wasted activations);
- The benefit per triggered beneficiary rises (earlier, better-targeted actions prevent more damage);
- The effective B:C ratio for the same investment base could approximately double relative to the current USD 7 baseline — approaching the USD 14–20 range for well-implemented programmes.

The USD 50–100 per person protected per pre-crisis cycle (WFP's approximate cost-per-beneficiary benchmark for anticipatory cash and in-kind support) remains the operational unit cost; τ's contribution is to ensure that a higher proportion of that per-person investment reaches genuine at-risk populations at actionable lead times.

11 Evidence and Translation Ladder

11.1 Phase 0: Benchmark and Validation (Months 0–12)

Goal: Establish that τ improves sub-seasonal and seasonal planning outputs on a selected set of high-priority tasks relative to current baselines.

Benchmark tasks:

- 30/60/90-day probabilistic rainfall and temperature anomaly envelopes, verified against WMO RCOF seasonal outlooks;
- Rainy-season onset and cessation date distributions, verified against ground-station records for 3–5 pilot countries;
- Dry-spell frequency and intensity during critical crop growth stages;
- District-level planting-window probability distributions, compared with extension-service planting recommendations;
- Drought-trigger lead-time performance: days of actionable advance warning relative to FFWC/IRI thresholds;
- Flood pre-positioning trigger performance: false-positive and false-negative rates relative to FFWC benchmarks.

KPIs: τ beats or matches benchmark for $\geq$70 percent of tasks; bounded error intervals are well-calibrated (coverage probability within 10 percent of stated confidence level); at least two national NMHS partners confirm operational relevance.

11.2 Phase 1: Climate-Service Insertion (Months 6–18)

Goal: Plug τ seasonal outputs into existing climate-service and seasonal-outlook pipelines in 3–5 pilot countries.

Insertion points:

- National meteorological and agricultural advisory systems;
- WMO RCOF and NCOF processes as a supplementary technical input;
- PICSA-style participatory climate-service programmes;
- Agricultural extension service seasonal advisory packages;
- District resilience dashboards.

KPIs: τ seasonal outlooks are officially incorporated into national agrometeorological advisory production in at least 2 countries; extension services report improved advisory specificity for planting-window guidance; humanitarian partner confirmation that τ inputs improve forward seasonal risk communication.

11.3 Phase 2: Anticipatory-Action Integration (Months 12–30)

Goal: Integrate τ seasonal outputs into forecast-based financing, early-action triggers, and shock-responsive social protection in 3–5 WFP anticipatory-action programme countries.

Natural partners: WFP country operations; FAO early warning/early action systems; national disaster-risk agencies; cash-transfer and rural-support systems; agricultural insurance and disaster-risk finance actors.

KPIs: Trigger recalibration documented in at least 3 WFP country programmes; false-activation rate reduced by $\geq$20 percent without increase in missed-activation rate; actionable lead time extended by $\geq$15 days in at least 2 programmes; cost-per-correctly-triggered beneficiary reduced.

11.4 Phase 3: Food-System Resilience Twin (Months 24–48)

Goal: Move from advisory fragments to a joined-up planning twin integrating climate, production, logistics, reserves, and social-protection signals.

Core layers: Seasonal climate outlook; production-risk envelope; water and pasture stress trajectory; storage and reserve position; social-protection readiness indicator; procurement and market-continuity status; emergency logistics pre-positioning status.

KPIs: Pilot national food-system resilience dashboard operational in at least 1 country, with all core layers updated at least monthly; documented use in at least one national reserve or social-protection decision; external evaluation confirms decision-grade product quality.

11.5 Phase 4: National and Regional Policy Integration (Months 36+)

Goal: Embed τ planning outputs in country investment planning and agrifood transformation processes.

Natural channels: World Bank FS-CAPs and climate-smart agriculture investment plans; national food-security dashboards and IPC coordination; strategic reserve modernization processes; public risk-finance systems; SADC/CILSS/IGAD regional food-security coordination; WMO RCOF institutionalization.

KPIs: τ seasonal planning product referenced in at least 2 FS-CAP processes; at least 1 regional food-security coordination body (SADC, CILSS, or IGAD) adopts τ seasonal inputs as part of its standard analytical package; GCF or World Bank funding authorized for regional expansion.

12 Stakeholder Map and Change Management

12.1 Primary institutional adopters

National meteorological and hydrological services (NMHSs) are the first-tier adopters. They are the institutional producers of seasonal climate services in their countries; τ integration means improving the quality of what NMHSs produce and distribute. NMHSs are typically under-resourced for major infrastructure change but highly motivated to improve seasonal forecast skill, which is a core measure of their professional standing.

WFP country operations and the Anticipatory Action team are the most operationally urgent adopters. They have existing trigger protocols, existing financing, and existing political relationships with host governments. τ integration requires trigger recalibration — a technically demanding but procedurally feasible change — and is unlikely to require major institutional restructuring.

Ministries of agriculture are the medium-term integrators. They receive seasonal climate advisories from NMHSs and use them for crop-campaign and reserve-policy guidance. Improved advisory quality reduces the political risk of early action and strengthens the justification for proactive reserve and safety-net decisions.

Social-protection ministries are the downstream beneficiaries of better trigger calibration. Their primary interest is in trigger reliability — they need to defend scale-up decisions to finance ministries. A physically grounded bounded-error trigger is more institutionally defensible than an ensemble-probability threshold.

12.2 Change management dynamics

The most significant institutional resistance is likely to come not from hostility to τ but from the perceived risk of replacing verified, institutionally trusted forecast products with novel ones. The change-management strategy should position τ as a complementary enhancement to existing products — running in parallel with IRI, ECMWF, and NMHS outputs, building a verification record, and demonstrating improvement incrementally — before proposing operational trigger recalibration.

The WMO RCOF process provides a natural institutional entry point: RCOFs already synthesize multiple forecast inputs and are designed to produce consensus seasonal outlooks. Introducing τ as an additional input into RCOF deliberations — rather than as a replacement for existing inputs — minimizes disruption and allows the value of τ's bounded-error outputs to be assessed against a well-understood institutional benchmark.

13 Gender, Equity, and Labor Dimensions

13.1 Women as primary food-system planners and shock absorbers

In most of the food-insecure systems targeted by this dossier, women bear disproportionate responsibility for household food-access management, water collection, and children's nutrition. They are also disproportionately exposed to the consequences of planning failure: when reserves run out, when planting is mistimed, or when anticipatory cash transfers arrive late, women's time burdens increase and their nutrition and income typically deteriorate first.

Better seasonal planning products therefore have asymmetric benefits for women — they disproportionately reduce the burden that falls on women when systems fail. Advisory and trigger systems should be explicitly designed to ensure that women's planting, storage, and fodder-management decisions are informed by the improved seasonal products, not only the decisions of male household heads or district-level extension contacts.

13.2 Pastoral and remote communities

Pastoral systems in the Sahel, the Horn of Africa, and Southern Africa are among the most underserved by current seasonal forecast and advisory products. Ground-station networks are sparse, and district-level probabilistic forecasts often have very low skill in pastoral zones. These communities face some of the highest climate-related food-security risk and the lowest access to planning-grade seasonal information.

Under τ, the bounded-error framework should in principle be applicable in data-sparse environments, since the bounds derive from structural physical constraints rather than from observational density. This makes τ potentially more equitable in coverage than observation-dependent ensemble forecast systems. Confirming this equity property should be an explicit validation criterion in Phase 0.

13.3 Labor and input-market disruption

Better-timed seasonal planning has labor-market implications. Optimal planting-window guidance, when adopted at scale, shifts peak demand for agricultural labor and inputs in ways that can either improve wage stability (by spreading seasonal demand more evenly) or disrupt existing informal labor-market arrangements. Social impact assessments should accompany any large-scale adoption of τ-based planting-window advisory systems to ensure that smallholders and agricultural laborers in the adoption zone do not face unforeseen labor-market disruption.

14 Benchmark Suite and Success Metrics

14.1 Forecast and planning skill

- 30/60/90-day probabilistic skill scores (Brier Skill Score or equivalent) for rainfall-anomaly, temperature-anomaly, and soil-moisture-anomaly forecasts, benchmarked against IRI seasonal outlooks and NMHS products;
- Rainy-season onset and cessation date error (root mean square error and mean absolute error against observed station records), benchmarked against WMO RCOF consensus outlooks;
- Drought-trigger lead-time performance: days of actionable advance warning before IPC Phase 3+ threshold is crossed, benchmarked against FEWS NET/IRI trigger timing;
- Flood pre-positioning trigger performance: false-positive and false-negative activation rates, benchmarked against FFWC threshold-based protocols.

14.2 Actionability and decision-quality metrics

- Proportion of τ seasonal outlooks that generate at least one concrete district-level advisory action in the partner country system;
- Advisory timeliness for seed, input, fodder, and reserve decisions: days in advance of the decision window that the advisory is issued;
- Percentage of funded anticipatory triggers that are activated on time relative to the realized shock envelope;

- Social-protection trigger alignment: proportion of τ-triggered activations confirmed as genuinely high-stress by ex-post IPC assessment.

14.3 Public-good outcomes

- Avoided crop and livestock losses, measured against counterfactual non-anticipatory-action scenarios in matched communities;
- Avoided distress sales: proportion of households reporting no forced asset sales during shock period among τ-triggered anticipatory-action beneficiaries;
- Households reached before shock peak: proportion of anticipatory-action beneficiaries receiving transfers or inputs before the IPC Phase 3 trigger threshold is crossed;
- Reserve-stock availability without fiscal overbuild: ratio of reserve stock utilized to reserve stock procured in τ-informed reserve-management pilots.

14.4 Equity and access metrics

- Smallholder coverage: proportion of households below 5 hectares among advisory recipients;
- Gender coverage: proportion of women as primary advisory recipients and primary cash-transfer recipients;
- Last-mile reach: proportion of pastoral and remote-community households with access to τ-informed seasonal advisories;
- Language and format accessibility: proportion of advisory outputs available in local languages and in low-connectivity (radio, SMS) formats.

15 Governance Guardrails

Deploying a physics-grade planning twin in food-system resilience contexts carries specific governance risks that must be addressed explicitly in implementation design.

1. **Protect local agency.** A stronger seasonal planning twin should support and augment local decision-making — by farmers, pastoralists, cooperatives, and local governments — not replace it with remote model authority. Advisory products must be designed to inform rather than prescribe; extension services and local institutions must retain the interpretive and decision authority that reflects local knowledge, local priorities, and local context.
2. **Keep uncertainty visible.** Even a stronger τ planning system must expose its uncertainty bounds and trigger logic transparently. Products that obscure residual uncertainty create false confidence that damages institutional trust when real uncertainty materializes. All τ planning products should display calibrated uncertainty intervals alongside point estimates.
3. **Validate before operationalizing triggers.** No anticipatory-action trigger protocol should be recalibrated using τ inputs before a minimum Phase 0 verification period demonstrates well-calibrated bounded-error performance in the relevant climate context. Premature trigger recalibration that degrades hit rates is worse than maintaining current baselines.
4. **Respect co-production norms.** WMO and WFP climate-service work consistently emphasizes co-production with users — PICSA, RCOF consultation, national forums — as essential to the uptake and trust-building that makes advisory products actionable. τ integration must operate within co-production frameworks, not bypass them.
5. **Design for smallholders and pastoralists, not only for ministries.** The highest social value in this application likely comes when improved seasonal planning reaches vulnerable producers,

not only national dashboards. Product design and dissemination channels must include extension services, radio, and SMS formats accessible to non-digitally-connected communities.

6. **Keep public interventions disciplined.** The World Bank/WFP/FAO grain-reserve guidance provides a useful reference: resilience tools can become costly, market-distorting, and politically difficult to reverse if objectives are unclear or if reserves are sized against worst-case scenarios rather than scenario-informed optimal profiles. τ should sharpen the calibration of public interventions, not provide cover for indiscriminate action.
7. **Avoid conflict with existing institutional trust relationships.** NMHSs, WFP country operations, and national food-security agencies have built trust relationships with governments over many years. τ integration should position itself as a technical improvement to inputs that NMHSs and WFP continue to own and deliver, not as a competing institutional product.
8. **Monitor for basis risk and distributional effects.** In parametric insurance and social-protection trigger applications, τ-based improvements should be continuously monitored for basis risk — the divergence between the index trigger and actual household-level conditions. If basis risk is higher in pastoral zones, remote areas, or among women-headed households, design corrections should be implemented before scaling.

16 SDG Mapping and Bottom Line

16.1 SDG alignment

SDG 2 — Zero Hunger. Paper 4's core application — improving seasonal planning and anticipatory action to reduce acute hunger — is a direct contribution to SDG 2's targets on ending hunger, achieving food security, and promoting sustainable food systems. Improved seasonal forecast quality and better-calibrated anticipatory-action triggers directly address Target 2.1 (universal food access), Target 2.4 (resilient agricultural practices), and Target 2.c (food commodity market volatility reduction).

SDG 1 — No Poverty. Distress sales, herd depletion, and forced coping strategies during agricultural shocks are among the primary mechanisms by which climate events push households below the poverty line. Better anticipatory action that prevents these events directly contributes to Target 1.5 (reduced exposure and vulnerability to extreme events).

SDG 13 — Climate Action. Early-warning system improvement and the integration of climate risk into national planning are explicit targets under SDG 13 (Targets 13.1 and 13.3). τ integration into national climate services and FS-CAPs is a direct implementation pathway for these targets.

SDG 17 — Partnerships for the Goals. The multi-institutional architecture required for τ deployment — NMHSs, WFP, FAO, World Bank, national governments, and regional bodies — directly embodies the partnership and means-of-implementation dimension of SDG 17, particularly in the science-policy-practice bridge function.

16.2 Bottom line

Paper 4 asks the most consequential practical question in the τ agriculture portfolio: can better seasonal physics become better food-system protection?

Under the τ assumption, the answer is yes — and the mechanism is not mysterious. The world has already built much of the institutional scaffold: agrometeorological services, seasonal climate forums, anticipatory-action programmes, climate-smart agriculture plans, food-security dashboards, and strategic reserve frameworks. What is still missing in many systems is a planning core that is physically stronger, more coherent across spatial scales, and more directly translatable into decision-grade action at the district level and seasonal horizon.

If τ can provide that core, this is not merely a better forecast story. It is a better protection story — one with a documented path from physics to planning to population-level outcomes, a viable finance architecture, a scalable deployment ladder, and a direct connection to the most urgent public-good challenge in the agriculture portfolio: preventing the next generation of avoidable food crises.

17 References

1. FAO. (2025). *Disasters cost global agriculture $3.26 trillion over three decades, FAO report reveals.* Food and Agriculture Organization of the United Nations. https://www.fao.org/newsroom/detail/disasters-cost-global-agriculture--3.26-trillion-over-three-decades--fao-report-reveals/en
2. WMO. (2025). *Agricultural Meteorology.* World Meteorological Organization. https://wmo.int/site/knowledge-hub/programmes-and-initiatives/agricultural-meteorology
3. WMO. (2025). *Regional Climate Outlook Forums and Regional Climate Forums.* World Meteorological Organization. https://wmo.int/activities/csis/regional-climate-outlook-forums-and-regional-climate-forums
4. WMO / FAO. (2024). *Empowering Farmers: The Role of Agrometeorological Services in Sustainable Agriculture.* World Meteorological Organization. https://wmo.int/media/update/empowering-farmers-role-of-agrometeorological-services-sustainable-agriculture
5. WMO. (2024). *Sub-seasonal to Seasonal Applications for Agriculture and Environment Project.* World Meteorological Organization. https://community.wmo.int/governance/sub-seasonal-seasonal-applications-agriculture-and-environment-project
6. WFP. (2025). *Anticipatory Action for Climate Shocks.* World Food Programme. https://www.wfp.org/anticipatory-actions
7. WFP. (2025). *Climate and Resilience.* World Food Programme. https://www.wfp.org/climate-and-resilience
8. WFP / FSIN / Global Network Against Food Crises. (2025). *Global Report on Food Crises 2025.* World Food Programme. https://www.wfp.org/publications/global-report-food-crises-grfc
9. World Bank. (2025). *Climate-Smart Agriculture and Food System Climate Action Plans (FS-CAPs).* The World Bank. https://www.worldbank.org/en/topic/climate-smart-agriculture
10. World Bank. (2025). *Food Security Update No. 117 — June 2025.* The World Bank. https://thedocs.worldbank.org/en/doc/40ebbf38f5a6b68bfc11e5273e1405d4-0090012022/related/Food-Security-Update-117-June-13-2025-2.pdf
11. World Bank, WFP, and FAO. (2025). *Strengthening Strategic Grain Reserves to Enhance Food Security.* World Bank. https://documents1.worldbank.org/curated/en/099042625211562573/pdf/P504545-488431b2-0565-40f9-852c-e8db32d22559.pdf
12. WFP. (2025). *Climate Services.* World Food Programme. https://www.wfp.org/climate-services
13. FEWS NET. (2025). *Famine Early Warning Systems Network.* USAID. https://fews.net

14. FAO GIEWS. (2025). *Global Information and Early Warning System.* Food and Agriculture Organization of the United Nations. https://www.fao.org/giews/en/
15. WFP VAM. (2025). *Vulnerability Analysis and Mapping.* World Food Programme. https://www.wfp.org/vulnerability-analysis-and-mapping
16. IRI Columbia. (2025). *International Research Institute for Climate and Society.* Columbia University. https://iri.columbia.edu
17. JRC / EU Copernicus. (2025). *MARS Agricultural Monitoring Service.* European Commission Joint Research Centre. https://joint-research-centre.ec.europa.eu/monitoring-agricultural-resources-mars_en
18. OCHA CERF. (2025). *Anticipatory Action — CERF Anticipatory Action Pillar.* United Nations Office for the Coordination of Humanitarian Affairs. https://cerf.un.org/anticipatory-action
19. GCF. (2025). *Climate Resilience and National Adaptation Plans.* Green Climate Fund. https://www.greenclimate.fund/topics/adaptation
20. SADC. (2024). *Declaration of Drought Disaster — Southern African Development Community Regional Humanitarian Emergency.* Southern African Development Community Secretariat. https://www.sadc.int
21. WFP Southern Africa. (2024). *El Niño 2023–24 Regional Response — Southern Africa.* World Food Programme Regional Bureau for Southern Africa.
22. Bangladesh FFWC. (2025). *Flood Forecasting and Warning Centre.* Bangladesh Water Development Board. https://www.ffwc.gov.bd
23. WMO / FAO ACREI. (2024). *Agricultural Climate Risk and Early Warning Integration — Greater Horn of Africa Programme.* World Meteorological Organization / Food and Agriculture Organization.
24. OCHA. (2025). *Southern Africa Drought Humanitarian Response, 2024.* United Nations Office for the Coordination of Humanitarian Affairs. https://www.unocha.org/southern-africa
25. START Network. (2025). *Anticipation Window.* START Network. https://startnetwork.org/start-anticipation-window
26. ECHO / DIPECHO. (2025). *Disaster Preparedness Programme.* European Commission Humanitarian Aid and Civil Protection. https://civil-protection-humanitarian-aid.ec.europa.eu/what/humanitarian-aid/disaster-preparedness_en
27. World Bank. (2025). *PROGREEN — Global Partnership on Sustainable and Resilient Food Systems.* The World Bank. https://www.worldbank.org/en/programs/progreen
28. WMO. (2024). *Greater Horn of Africa Climate Outlook Forum (GHACOF).* World Meteorological Organization. https://www.icpac.net/ghacof/
29. IGAD Climate Prediction and Applications Centre (ICPAC). (2025). *Seasonal Climate Outlook for Greater Horn of Africa.* https://www.icpac.net
30. CILSS. (2025). *Comité permanent Inter-Etats de Lutte contre la Sécheresse dans le Sahel — Food Security and Agrometeorological Products.* https://www.cilss.int
31. IPC Global Support Unit. (2025). *Integrated Food Security Phase Classification — IPC Acute Food Insecurity.* https://www.ipcinfo.org
32. WMO. (2024). *WMO Guidelines on Multi-Hazard Impact-Based Forecast and Warning Services.* World Meteorological Organization.

33. FAO. (2024). *State of Food and Agriculture 2024: Financing to End Hunger and Malnutrition.* Food and Agriculture Organization of the United Nations.
34. AGRHYMET. (2025). *Regional Centre for Agrometeorology and Hydrology — Sahel Seasonal Outlook Products.* `https://www.agrhymet.ne`
35. Alfieri, L. et al. (2023). *Advancing Global Flood Early Warning with Improved Sub-seasonal Forecasts.* Nature Hazards and Earth System Sciences.
36. Barnston, A. G., et al. (2023). *Skill and Reliability of ENSO-Based Seasonal Forecasts for Agricultural Planning.* Bulletin of the American Meteorological Society.
37. Coughlan de Perez, E., et al. (2022). *Forecast-Based Financing: From Research to Humanitarian Practice.* Natural Hazards and Earth System Sciences.
38. Hess, U., & Hazell, P. (2016). *Innovations and Emerging Trends in Agricultural Insurance.* Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ).
39. IPCC. (2022). *Climate Change 2022: Impacts, Adaptation, and Vulnerability — Chapter 5: Food, Fibre, and Other Ecosystem Products.* Intergovernmental Panel on Climate Change.
40. Tall, A., et al. (2014). *Scaling Up Climate Services for Farmers in Africa and South Asia.* CGIAR Research Program on Climate Change, Agriculture, and Food Security.

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 agro-climate-resilience twin could strengthen sub-seasonal to seasonal planning, anticipatory action, and food-system resilience at district, national, and regional scales. The v3 impact thesis is conditional: a Tau-grade seasonal climate-food-system resilience planning 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 agro-climate-resilience twin could strengthen sub-seasonal to seasonal planning, anticipatory action, and food-system resilience at district, national, and regional scales. 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

- **FAO**, State of the World's Land and Water Resources for Food and Agriculture 2025 [1]: freshwater, land, irrigation, and food-system baseline.
- **World Bank Group**, Transforming Lives Through Climate-resilient Irrigation [8]: climate-resilient irrigation and productivity baseline.
- **FAO**, WaPOR: Remote Sensing for Water Productivity [3]: water productivity and evapotranspiration data baseline.
- **FAO**, AquaCrop Crop-Water Productivity Model [2]: crop-water model comparison baseline.
- **WMO**, Agricultural Meteorology Programme [7]: public agrometeorological service context.
- **IFAD**, Climate Change and Resilience [5]: smallholder adaptation and rural resilience context.

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 FAO, IFAD, 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 seasonal climate-food-system resilience planning 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: **seasonal climate-food-system resilience planning twin**.

First benchmarkable test: seasonal crop, food-security, and disruption-risk scenarios against FEWS NET, national crop forecasts, and disaster-loss records.

- 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 Seasonal Planning, Disaster Anticipation, and Food-System Resilience without operational authority.	medium
Realistic	A reviewed prototype strengthens several public-sector workflows for Seasonal Planning, Disaster Anticipation, and Food-System Resilience after benchmark comparison with incumbent systems.	medium-low
Optimistic	A reusable public-good intelligence layer becomes plausible for Seasonal Planning, Disaster Anticipation, and Food-System 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	4/5	The evidence base is anchored in public institutions, official monitoring systems, or established scientific reviews.
Measurability	4/5	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

seasonal planning exercise with a food-security agency and national agriculture ministry

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	seasonal crop, food-security, and disruption-risk scenarios against FEWS NET, national crop forecasts, and disaster-loss records	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] FAO. State of the world's land and water resources for food and agriculture 2025. https://www.fao.org/land-water/solaw2025/en/, 2025. freshwater, land, irrigation, and food-system baseline.

[2] FAO. Aquacrop crop-water productivity model. https://www.fao.org/aquacrop/en/, 2026. crop-water model comparison baseline.

[3] FAO. Wapor: Remote sensing for water productivity. https://www.fao.org/in-action/remote-sensing-for-water-productivity/en/, 2026. water productivity and evapotranspiration data baseline.

[4] Thorsten Fuchs and Anna-Sophie Fuchs. Tau for seasonal planning, disaster anticipation, and food-system resilience. https://panta-rhei.site/impact/papers/seasonal-planning-disaster-anticipation-food-system-resilience/, 2026. Current public full-text source for dossier seasonal-planning-disaster-anticipation-food-system-resilience.

[5] IFAD. Climate change and resilience. https://www.ifad.org/en/climate-and-environment, 2026. smallholder adaptation and rural resilience context.

[6] Panta Rhei Research Program. Public-good briefing landing page. https://panta-rhei.site/publications/research-briefings/public-good/seasonal-planning-disaster-anticipation-food-system-resilience/, 2026.

[7] WMO. Agricultural meteorology programme. https://wmo.int/activities/agricultural-meteorology, 2026. public agrometeorological service context.

[8] World Bank Group. Transforming lives through climate-resilient irrigation. https://shorthand.worldbankgroup.org/transforming-lives-through-climate-resilient-irrigation/, 2026. climate-resilient irrigation and productivity baseline.

Panta Rhei Research Program

Public-Good Impact Dossier

Tau for Seasonal Planning, Disaster Anticipation, and Food-System Resilience

Dossier ID: PGID-AGRI-04 Portfolio: Agriculture 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