Panta Rhei
Research Program

Research Briefings · Public-Good Impact Dossiers

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

Tau for Pest, Disease, and Livestock-Stress Early Warning

Conditional public-good pathway for Pest, Disease, and Livestock-Stress Early Warning

Public-Good Impact Dossier

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

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

Thorsten Fuchs • Anna-Sophie Fuchs

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

Safe harbor / release discipline

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

Briefing identity and source routes

Dossier ID	PGID-AGRI-03
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/pest-disease-livestock-stress-early-warning/`
Landing route	`https://panta-rhei.site/publications/research-briefings/public-good/pest-disease-livestock-stress-early-warning/`

Release status

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

What this dossier claims

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

What this dossier does not claim

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

Source fidelity and legal disclaimer

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

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

Contents

1 Executive Summary 4

2 Why This Matters Now 5

3 Scope and Reader Orientation 6

4 The Opportunity Baseline 6

5 Working τ Assumptions 8

6 What Changes with a Law-Faithful Twin 8

7 Competitive and Incumbent Landscape 10

8 Structured Opportunity Map 12

9 Geographic Case Studies 14

10 Finance, ROI, and Climate-Finance Eligibility 17

11 Evidence and Translation Ladder 18

12 Stakeholder Map and Change Management 19

13 Gender, Equity, and Labor Dimensions 20

14 Benchmark Suite and Success Metrics 21

15 Governance Guardrails 22

16 SDG Mapping and Bottom Line 23

17 References 23

18 Dossier accountability addendum 26

1 Executive Summary

This dossier examines what a physically faithful, law-grounded discrete twin of weather–pathogen–vector–host dynamics — built on the τ framework — could deliver for plant-pest and crop-disease early warning, transboundary animal-disease intelligence, livestock heat-stress management, and One Health biosecurity integration.

The underlying burden is large, well-documented, and institutionally acknowledged. FAO reports that up to 40 percent of global food crops are lost annually to plant pests, generating more than USD 220 billion in annual trade losses and at least USD 70 billion in invasive-pest economic losses. Annual livestock losses from transboundary animal diseases range from USD 48 to 330 billion, with aquaculture adding USD 10 billion per year. Extreme heat threatens the livelihoods of 1.23 billion people in agrifood systems; agricultural workers are 35 times more likely to die from occupational heat exposure than workers in other sectors. The five key findings that give structure to this dossier are as follows.

Key Finding 1. Plant-health and animal-health early warning already works through mature international systems — FAO DLIS, FAMEWS, EMPRES-i+, and WOAH WAHIS — but all current systems depend on loosely coupled stacks: weather forecasting in one module, field surveillance in another, ecology and movement models in a third, and local decision rules that are often heuristic and coarsely calibrated. A τ-enabled twin would replace this patchwork with a single coherent physical substrate in which weather, vectors, pathogens, and hosts operate under the same governing laws.

Key Finding 2. Earlier warnings are only valuable if they are spatially and operationally specific enough to trigger targeted action. τ-enabled outputs would improve not only when a warning is generated but where it applies, how rapidly confidence changes, and what intervention window remains open — reducing both missed outbreaks and costly overreaction in the form of blanket spraying, indiscriminate culling, and over-broad movement restrictions.

Key Finding 3. The competitive and incumbent landscape includes strong public systems (FAO DLIS, FAMEWS, EMPRES-i+, WOAH WAHIS) and commercial platforms (Cropin, Plantix, SatAgro, CABI Plantwise). No incumbent currently integrates plant-health, animal-health, vector ecology, and livestock heat stress in a single law-faithful forecasting engine — that integration gap is τ's primary point of differentiation.

Key Finding 4. Geographic pilots with the clearest near-term leverage are: the East Africa desert locust corridor (where FAO's DLIS response in 2020 cost USD 120 million and still faced localization and timing gaps), West and Central Africa fall armyworm zones (FAO-estimated USD 13 billion per year in potential losses), and Kenya/East Africa Rift Valley Fever corridors (where existing FAO tools were not fully operationalized for the 2023 outbreak).

Key Finding 5. The finance ecosystem for this domain is substantial and growing: the World Bank's One Health Platform has committed approximately USD 2 billion across IDA and IBRD lending; GCF's anticipatory-action window is open to national-level biosecurity applications; FAO's Emergency and Resilience Trust Fund supports early-action deployments; and WOAH's Global Health Emergency Fund specifically targets animal-disease preparedness. Cost scenarios of USD 5–12 million per country and USD 30–60 million for a regional East Africa/Sahel corridor platform are consistent with this funding landscape.

This is a yellow paper: assumption-led, translation-oriented, and public-good framed. It does not claim that the agricultural, veterinary, or scientific communities have accepted the τ assumptions. It asks what would follow — institutionally, financially, and operationally — if those assumptions hold well enough to matter.

2 Why This Matters Now

2.1 The climate-disease convergence is accelerating

The relationship between climate variability and biological risk is not linear and is not simply a matter of warming averages. It operates through threshold effects: a single rainfall pulse that arrives outside its historical window can trigger locust breeding in an area previously considered unsuitable; a three-day temperature anomaly can push a vector species into a new elevation band; a soil-moisture deficit combined with thermal stress can suppress immune function in livestock before any infectious agent is present. Climate change is systematically widening the geographic range, extending the seasonal window, and compressing the warning-to-impact timeline for a large class of biological threats.

FAO's 2025 plant-health early-warning assessment frames this directly: climate shifts are expanding pest ranges into previously unaffected territories, enabling year-round survival of species that were previously controlled by seasonal cold, and increasing the volatility of weather sequences on which current forecasting models were calibrated. The same assessment notes that early warning systems for plant health protect not only crops and food security, but also human health, biodiversity, and the planetary boundaries associated with chemical pesticide use.

The WMO/FAO joint 2025 report on extreme heat and agriculture adds a parallel dimension: heat events that once occurred once per decade are now occurring three to five times per decade in major agricultural zones, and the biological responses — in crops, in livestock, and in disease vectors — are non-linear. This means the error cost of late or imprecise warnings is rising even as the frequency of events requiring a warning is also rising.

2.2 The One Health imperative is no longer optional

The COVID-19 pandemic made One Health — the integrated management of human, animal, and environmental health — a mainstream policy priority. But the operational implementation of One Health in agriculture remains fragmented. Plant-health surveillance systems do not communicate in real time with animal-health surveillance systems. Weather forecast services do not feed directly into vector-ecology risk engines. National plant-protection organizations and national veterinary authorities operate under separate legal mandates, separate data architectures, and separate funding streams.

FAO's latest framing is unambiguous: the boundaries between plant health, animal health, human health, and environmental health are artificial from a risk-management standpoint. A drought that stresses crops in the same corridor where Aedes mosquitoes breed seasonally can simultaneously increase fall armyworm pressure, trigger Rift Valley Fever vector expansion, and reduce livestock feed security — all within the same three-week weather window. Managing that compound event through separate institutional pipelines is both less effective and more expensive than managing it through a unified biosecurity layer.

2.3 Migratory pests are a stress test the current system has already failed

The 2020–2021 desert locust outbreak was the most severe in 25 years for East Africa and 70 years for the Horn. FAO's emergency response cost approximately USD 120 million in 2020 alone. FAO's DLIS provided operational six-week outlooks and was critical to mobilizing the response — but ground response was delayed in several areas due to conflict, access constraints, and imprecision in swarm-location forecasting. Aerial treatment operations covered large areas where swarms were not present, while swarms in neighboring corridors were missed due to vegetation and rainfall gaps in the forecast model.

Fall armyworm, which first appeared in Africa in 2016, spread to 44 countries within two years. FAO estimates potential losses of USD 13 billion per year across affected maize-growing areas. Maize constitutes 60 to 70 percent of caloric intake in sub-Saharan Africa; the welfare stakes of an improved early warning layer are not marginal.

2.4 Livestock heat stress is transitioning from background factor to acute risk

WMO reports that 470 billion labor hours were lost globally due to extreme heat in 2021. Agricultural workers face a mortality risk from occupational heat exposure 35 times higher than workers in other sectors. For livestock themselves, heat stress directly affects feed intake, fertility, milk and egg production, transport safety, disease susceptibility, and mortality. Yet most operational heat-advisory systems available to farmers today offer only coarse regional temperature alerts without the species-specific, production-stage-specific, and microclimate-aware granularity needed to trigger precise protective actions.

The institutional demand signal is already present. The regulatory and data infrastructure already exists. What is missing is a physics-faithful engine that can translate weather and ecological state into decision-relevant biological risk at the farm and district level — consistently, without arbitrary parameter tuning, and across all three biosecurity channels simultaneously.

3 Scope and Reader Orientation

This is Paper 3 of 5 in the Panta Rhei Impact Agriculture Portfolio. It focuses on the early-warning and anticipatory-action layer for biological threats to agricultural systems: plant pests, crop diseases, transboundary animal diseases, vector-borne diseases, and livestock heat stress.

In scope for this paper: - Plant-pest and crop-disease early warning (migratory pests, weather-sensitive crop disease windows) - Transboundary animal-disease intelligence and vector-borne disease forecasting - Livestock heat-stress, welfare, and productivity windows - Integrated weather–ecology–host decision support - One Health fusion: shared causal engine for plant, animal, human, and environmental health risks

Explicitly out of scope for this paper: - Paper 1 covers operational agro-weather intelligence for routine field decisions. - Paper 2 covers climate-smart irrigation, soil moisture, and water productivity. - Paper 4 covers seasonal planning, disaster anticipation, and food-system resilience. - Paper 5 covers crop biology, breeding, photosynthesis engineering, and targeted gene design.

This paper adopts a yellow-paper stance throughout: it assumes, for planning purposes, that the τ framework's claims relevant to this domain are sound, and asks what practical consequences would follow if those claims were integrated into existing plant-health, animal-health, and farm-advisory systems. Readers who want to evaluate the theoretical foundations of τ are referred to the Panta Rhei series itself. This paper is addressed primarily to decision-makers in ministries of agriculture, plant-protection organizations, veterinary services, national meteorological and hydrological services, food-safety authorities, farmer cooperatives, livestock boards, extension systems, humanitarian actors, One Health platforms, and climate/agriculture funders.

4 The Opportunity Baseline

4.1 Plant health: a USD 220+ billion annual burden

FAO's 2025 plant-health assessment provides the clearest institutional baseline. Up to 40 percent of global food crops are lost to plant pests annually. Annual trade losses due to plant pests in agricultural products exceed USD 220 billion. Invasive pests cause at least USD 70 billion in additional annual global economic losses. These figures represent a burden that is not merely persistent but is growing: climate change is expanding the geographic range and seasonal window of many of the most damaging species, while global trade and travel networks continue to accelerate the movement of pests into new territories.

The public-good framing matters here. The FAO assessment notes that plant diseases and pests are not only an agricultural productivity issue. They are a human nutrition issue, a biodiversity issue, a chemical-use issue, and a food-security issue for the most vulnerable populations. When early warning fails and a major outbreak occurs, the harm cascades: yield losses, emergency pesticide applications with collateral ecological damage, trade disruptions, food price spikes, and livelihood shocks that can take years to recover from.

4.2 Animal health: a USD 48–330 billion annual burden

FAO's November 2025 transboundary animal disease report sets out the scale of the livestock-health burden with similar clarity. The global farmed-animal sector is valued at USD 1.6 to 3.3 trillion. Annual livestock losses from transboundary animal diseases are estimated at USD 48 to 330 billion. Aquaculture adds USD 10 billion in annual disease-related losses. Livestock underpin the livelihoods of 1.9 billion people, and outbreaks can erase years of development gains within days.

The width of the USD 48–330 billion range is itself informative: it reflects genuine uncertainty about the true scope of losses, not a deliberate hedge. This uncertainty is partly a data-quality problem and partly a structural property of the system — losses materialize through multiple channels (mortality, production decline, trade access, emergency response costs, secondary human health impacts) that are tracked inconsistently across countries and disease systems.

4.3 Livestock heat stress: 1.23 billion livelihoods at risk

FAO and WMO's 2025 joint assessment on extreme heat and agriculture reports that 1.23 billion people in agrifood systems have their livelihoods threatened by extreme heat. Agricultural workers are 35 times more likely to die from occupational heat exposure than workers in other sectors. In 2021 alone, 470 billion labor hours were lost globally due to extreme heat, the vast majority of them in the agricultural sector in low- and middle-income countries.

For livestock systems specifically, heat stress affects dairy cattle, poultry, pigs, small ruminants, and draught animals through a set of well-characterized physiological mechanisms: reduced feed intake, elevated maintenance energy requirements, suppressed reproductive performance, reduced immune competence, and — in severe cases — acute mortality. The economic losses from these effects are difficult to aggregate globally but are locally severe: a single heat event lasting three to five days can reduce milk production in an unprotected dairy system by 10 to 25 percent, and repeated heat exposure over a season can affect fertility outcomes for an entire reproductive cycle.

4.4 The demand signal is institutionally mature

These are not problems searching for institutions. FAO's DLIS has operated continuously for decades. FAMEWS has provided a real-time global platform since fall armyworm first reached Africa. EMPRES-i+ maps hundreds of thousands of disease records across 190 countries. WOAH's WAHIS serves as the official early-warning and monitoring backbone for international animal-disease governance. The infrastructure, the mandates, and the operational experience are all present. What the current systems cannot yet fully provide is a unified, law-faithful, real-time coupling between weather dynamics and the biological risk engines they feed.

5 Working τ Assumptions

This dossier adopts four categories of τ assumptions for planning purposes. These are stated explicitly so that readers can evaluate the conditional nature of the scenarios that follow.

5.1 Physics-side assumptions

The τ framework provides a discrete, constructive, countable, bounded-error substrate for weather dynamics: atmospheric moisture transport, rainfall pulses, thermal fields, surface water, vegetation state, and microclimatic conditions relevant to pest emergence, pathogen viability, vector breeding, and livestock heat exposure. This substrate is coarse-grained but law-faithful: it does not drift arbitrarily as refinement depth increases, and it can be coupled to field observations without losing structural coherence. Under this assumption, τ can translate atmospheric-state sequences into operational field conditions — outbreak risk windows, habitat suitability maps, vector emergence probabilities, grazing stress indices, and thermal danger periods — without requiring arbitrary local parameter tuning.

5.2 Biology-side assumptions

The τ biological framework is rich enough to represent host–pathogen–vector interactions at a decision-relevant forecasting level. This does not require perfect predictability of every disease or biological pathway. It requires instead that τ can capture the physics-limited and ecology-limited components of the problem — the parts where weather sequences, habitat conditions, and transport dynamics drive risk — and that this physics layer is improved enough to deliver materially better operational outputs than current stacked models. τ can map ecological state to decision-relevant risk (not merely generic correlation scores) and can ingest surveillance data, field observations, and sensor inputs while maintaining coherent causal structure.

5.3 Operational assumptions

τ outputs are designed to feed existing operational systems rather than replace them. This includes locust-monitoring platforms, plant-protection advisory services, veterinary surveillance systems, livestock management channels, and humanitarian early-action protocols. τ outputs are structured as decision-relevant products: where to scout, where to place traps, when to treat or withhold treatment, when to vaccinate or intensify surveillance, when to restrict animal movement, and when to change watering, shading, or grazing schedules.

5.4 What this dossier does not assume

This paper does not assume perfect predictability of biological outcomes, linear scalability of τ outputs across all geographies and species, or that institutional adoption will be frictionless. It does not claim that the τ framework eliminates the need for field surveillance — on the contrary, a stronger physics engine should sharpen the targeting and timing of field surveillance, not substitute for it. It also does not claim that the benefits described in this paper are certain or assured. They are planning-inference scenarios conditional on the τ assumptions holding well enough to matter operationally.

6 What Changes with a Law-Faithful Twin

6.1 From a stacked architecture to a unified causal substrate

Today's early-warning systems reflect a familiar architectural split. Weather forecasting operates in one stack, managed primarily by national meteorological services and global NWP centers. Field surveillance — scouting reports, pheromone-trap data, satellite vegetation indices — operates in a second stack, often managed by extension services, plant-protection organizations, or veterinary surveillance networks. Ecology and movement models occupy a third layer, often implemented as empirical statistical models or expert-rule systems calibrated on historical outbreak data. Local decision rules sit atop all of these and are typically heuristic: spray when trap counts exceed threshold X, vaccinate when the RVF rainfall index exceeds Y.

The coupling between these stacks is loose. Weather forecasts feed risk indices that were derived from historical correlations. Field surveillance updates maps but does not always feed back into the weather-physics layer. Ecology models are often not re-calibrated when weather patterns shift. The result is a system where the gap between physical state and operational decision is wide enough to allow both costly delays and costly false alarms.

Under the strongest τ assumption, that split weakens fundamentally. The operational system would execute the same structural laws that the weather–ecology–host system itself obeys, at a certified coarse-grained resolution. Weather, vectors, pathogens, and host stress would share a single substrate, and field observations would update a genuinely causal risk engine rather than a loosely assembled dashboard.

6.2 Earlier warnings that are operationally specific

Lead time improvements only reduce losses if the warning is specific enough to trigger targeted action. The desert locust example is instructive: FAO's DLIS provided six-week outlooks and these were critical to the 2020 response — but the spatial resolution of swarm-location forecasting was insufficient to prevent both missed areas and unnecessary interventions in areas where no swarms appeared. A τ-enabled twin would improve not only the temporal lead of a warning but its spatial precision, its confidence evolution over the forecast horizon, and its translation into specific scouting priorities and treatment areas.

For vector-borne animal diseases, the same dynamic applies. FAO's Rift Valley Fever decision-support tool can issue alerts one to two months before first infections are reported — but the 2023 Kenya outbreak revealed gaps in operationalization. τ would tighten the coupling between rainfall-triggered flooding events, Aedes mosquito breeding habitat probability, and livestock vaccination-campaign timing, making the interval between environmental trigger and protective action shorter and more precisely targeted.

6.3 Reduction in false positives and blanket interventions

A major hidden cost of weak early warning is not missed outbreaks but overreaction. Blanket pesticide applications cover areas where pest pressure is absent; emergency movement restrictions affect farms and corridors where disease risk is low; vaccination campaigns are timed to worst-case envelope rather than actual risk. These interventions carry costs — financial, ecological, and social — and they erode trust in early-warning systems among the farmers and herders whose cooperation is essential for ground-truthing.

A more faithful twin should lower this waste systematically. By improving the spatial and temporal specificity of risk estimates, τ would allow plant-protection organizations and veterinary authorities to concentrate resources where risk is highest, withdraw advisory intensity where it is low, and provide farmers with more actionable guidance that matches their operational decision windows.

6.4 One Health fusion: plant, animal, and heat warnings from one engine

The most structurally significant change would be the consolidation of plant-health, animal-health, and livestock heat-stress warning into a single causal engine. Under current institutional architecture, a compound event — a rainfall pulse that simultaneously triggers locust breeding, increases Aedes mosquito density, and reduces livestock thermal resilience through humidity — is processed through three separate operational channels that do not communicate in real time. The result is sequential, reactive, and resource-inefficient.

A τ early-warning twin capable of supporting all three channels from the same local weather–ecology substrate would enable genuinely integrated biosecurity response: a single national platform that generates coordinated plant-health, animal-health, and livestock-welfare advisories from the same physical state estimate.

6.5 Anticipatory action becomes more defensible

If the warning engine is more trustworthy, actions taken before losses occur become easier to justify politically and financially. Pre-positioning treatment materials, releasing humanitarian funding, initiating vaccination campaigns, or restricting movement ahead of a confirmed outbreak all carry institutional risk if the underlying warning proves wrong. A τ-grounded warning with certified error envelopes and transparent decision logic is more defensible — to legislators, to funders, and to affected communities — than a warning derived from loosely coupled empirical models. This matters especially for anticipatory financing instruments, parametric insurance triggers, and humanitarian early-action protocols.

7 Competitive and Incumbent Landscape

The early-warning space for plant pests, animal diseases, and livestock stress is populated by a mix of mature international public institutions, global advisory platforms, and commercial precision-agriculture services. No single incumbent currently provides a unified, physics-faithful, multi-hazard biosecurity twin. The gap that τ addresses is structural, not marginal.

7.1 FAO DLIS / FAMEWS / EMPRES-i+

Capabilities. FAO operates three distinct but complementary early-warning services. The Desert Locust Information Service (DLIS) is the world's oldest migratory-pest early-warning system, operating 24/7 with six-week locust outlooks, six-month precipitation predictions, field officer tools, and global coverage of all desert locust-affected regions. The Fall Armyworm Monitoring and Early Warning System (FAMEWS) provides real-time global and sub-national infestation maps fed by scouting and pheromone-trap data from national plant-protection organizations across Africa and Asia. EMPRES-i+ maps hundreds of thousands of animal-disease records across 190 countries and integrates with GLEWS+ for global disease intelligence, while the RVF decision-support tool can issue alerts one to two months ahead of first reported infections.

Geographic scope. Global, with strongest coverage in regions that have invested in national reporting infrastructure.

Differentiation from τ. FAO's systems are strong on institutional coverage, ground-truth reporting networks, and operational experience. They are weaker on sub-field precision in pest location forecasting, on real-time coupling between weather dynamics and biological risk engines, and on integration across plant-health, animal-health, and heat-stress channels. τ would add a law-faithful physical substrate that the current FAO stacks do not have, improving forecast precision and enabling genuine One Health fusion.

7.2 CABI Plantwise

Capabilities. CABI's Plantwise programme operates in approximately 90 countries, supporting a network of plant clinics staffed by trained plant doctors. The programme offers digital disease diagnostics, extension-linked advisory services, and a growing data platform for plant-health information. CABI also supports national plant-protection organizations with diagnostic capacity building and pest-risk analysis tools.

Geographic scope. Global, with strongest implementation in sub-Saharan Africa, South Asia, and parts of Latin America; primarily extension-facing at the local level.

Differentiation from τ. Plantwise is exceptional at the local diagnostic and extension layer — the human interface between a plant health problem and the farmer who needs to respond. It is not primarily a forecast-led early-warning system. It does not generate outbreak predictions from physics, does not couple weather dynamics to disease risk, and is not designed for proactive biosecurity intelligence ahead of pest arrival. τ would complement Plantwise by providing the upstream forecast layer that feeds Plantwise's extension network with targeted scouting priorities and advisory windows.

7.3 Cropin

Capabilities. Cropin is an Indian agri-data intelligence company operating across 22 million-plus acres globally. Its platform integrates satellite imagery, weather data, crop-monitoring analytics, and machine-learning models for disease forecasting, yield prediction, and supply-chain intelligence. Cropin operates commercially across South Asia, Southeast Asia, Africa, and Latin America, primarily serving agribusinesses, cooperatives, and government clients.

Geographic scope. Global presence, strongest in South Asia; commercial SaaS model.

Differentiation from τ. Cropin demonstrates that integrated satellite-weather-disease intelligence is commercially viable and scalable. However, its forecasting engine is machine-learning-based, calibrated on historical correlations, and is not a physics-faithful twin. Its models do not encode the governing laws of weather-pathogen-vector dynamics; they learn statistical associations that may generalize poorly to new geographies, new climate regimes, or compound events outside the training distribution. τ's differentiation is not better data integration but a fundamentally different forecasting substrate that is law-faithful rather than correlation-based.

7.4 PEAT (Plantix)

Capabilities. Plantix is a mobile disease diagnosis application developed by PEAT GmbH, with more than 10 million downloads, primarily across South and Southeast Asia. The application uses AI image recognition to identify crop diseases, pests, and nutrient deficiencies from smartphone photographs, and provides treatment recommendations. Plantix is arguably the most farmer-facing tool in this landscape, with high adoption in smallholder systems.

Geographic scope. South Asia dominant, expanding into Africa and Southeast Asia; farmer-direct model.

Differentiation from τ. Plantix is a reactive diagnostic tool: it helps farmers identify a problem after it is visually apparent in the field. It does not provide proactive early warning from physical forecasting; it cannot predict outbreak windows or movement fronts; and it is not integrated with weather dynamics or vector ecology. τ operates at a fundamentally different point in the warning timeline — before field symptoms appear, at the level of environmental suitability and outbreak probability.

7.5 SatAgro

Capabilities. SatAgro is a European satellite-based crop monitoring and precision agriculture platform that integrates NDVI, weather data, and field scouting for crop health assessment, variable-rate application mapping, and farm management support. Its primary focus is precision agriculture in Eastern Europe, with growing interest in other European markets.

Geographic scope. Eastern Europe primary; expanding European market; commercial.

Differentiation from τ. SatAgro provides high-quality farm-level crop monitoring for precision agriculture contexts. It does not integrate vector or pest ecology, does not provide biosecurity early warning, and is not designed for transboundary or migratory pest systems. τ addresses a fundamentally different problem space: regional and national biosecurity intelligence rather than farm-level precision agronomy.

7.6 WOAH WAHIS

Capabilities. The World Organisation for Animal Health's World Animal Health Information System is the official regulatory backbone for international animal-disease notification and early warning. WAHIS collects immediate notifications of listed and emerging animal diseases from member countries, publishes six-monthly reports, and serves as the authoritative data layer for disease-outbreak intelligence worldwide. WAHIS is embedded in international trade and sanitary regulations; its notifications have legal and market-access implications.

Geographic scope. Global; 182 member countries; official regulatory function.

Differentiation from τ. WOAH WAHIS has authoritative institutional status and an unmatched reporting network. Its primary function is official notification and surveillance monitoring rather than real-time forecast generation. WAHIS does not generate predictive risk maps from weather or vector-ecology models; it records and communicates confirmed and suspected events. τ would sit upstream of WAHIS, feeding risk probability into the decision-making processes that precede official notification and enabling earlier anticipatory action before a WAHIS-reportable event occurs.

8 Structured Opportunity Map

8.1 Migratory and transboundary crop pests

This is the clearest near-term public-value beachhead in this portfolio.

Migratory pest systems — desert locust, fall armyworm, brown planthopper, whitefly — share a common structural property: their outbreak dynamics are strongly weather-driven and their movement is governed by atmospheric transport. The decision to breed, swarm, migrate, or establish is coupled to rainfall pulses, wind fields, vegetation green-up, and temperature sequences in ways that are physically deterministic at the scale of days to weeks. This is precisely the class of problem where a τ weather–ecology twin offers the greatest incremental precision over existing correlation-based models.

Under τ, the operational gains would be: - Stronger habitat suitability forecasting linked to real-time rainfall and vegetation dynamics. - More precise swarm genesis probability maps that distinguish confirmed breeding zones from speculative risk areas. - Tighter outbreak-front tracking as populations move across corridors. - Earlier treatment-window recommendations that allow targeted aerial or ground application before swarms reach damaging density. - Reduced aerial spray waste through better prioritization of treatment zones. - Integration with livestock-feed emergency planning when pasture areas are at risk.

FAO DLIS has already demonstrated the operational architecture for this type of system. τ's contribution is to improve the physics layer that feeds that architecture.

8.2 Weather-sensitive crop disease windows

Many economically important crop diseases are not continuously present threats; they intensify, spread, or become economically significant only under specific sequences of weather conditions. Wheat rust requires particular temperature and humidity bands; late blight in potato is highly sensitive to leaf wetness duration and temperature; Fusarium head blight requires specific moisture windows during anthesis; coffee leaf rust is driven by rainfall and temperature combinations that are shifting as climate changes.

Current crop-disease advisory systems typically provide regional “disease-favorable condition” indicators built from empirical correlations between historical weather records and outbreak data. These indicators are useful but coarse: they cannot reliably distinguish a three-day window of favorable conditions that will produce an economically significant outbreak from one that will not, and they cannot localize risk to the field level without dense sensor networks.

Under τ, the operational gains would be: - Decision-relevant local windows: when to scout, when to spray, when not to spray, and how long the action window will remain open. - Field-specific risk maps based on local microclimate, topography, and crop stage, not only regional weather averages. - Reduced unnecessary fungicide applications — both a cost saving and an environmental and resistance-management benefit. - Improved timing of biological control applications where crop stage and weather window are critical to efficacy.

This opportunity is less centralized than locust or fall armyworm systems, but the aggregate burden is large and the operational value of improved precision is very high.

8.3 Vector-borne and transboundary animal disease

Vector-borne animal diseases — Rift Valley Fever, bluetongue, African horse sickness, and increasingly lumpy skin disease — are driven by the ecology and abundance of insect vectors, which in turn are driven by rainfall, temperature, humidity, and surface water. The link from weather to vector population to disease outbreak is causal and physically determined; it is not merely correlational.

FAO's existing RVF decision-support tool already demonstrates that this causal chain can be captured operationally: the tool issues alerts one to two months before first reported infections by tracking rainfall and vegetation conditions associated with Aedes mosquito breeding. τ would strengthen this chain by providing a more precise, continuous, law-faithful representation of the weather–vector–outbreak pathway, reducing both the false-alarm rate and the probability of missed early signals.

For non-vector-borne transboundary diseases — foot-and-mouth disease, highly pathogenic avian influenza, African swine fever — the weather coupling is indirect but still important: temperature affects virus survival in the environment, extreme events affect animal movement and congregation patterns, and habitat conditions affect the mixing rates of wild and domestic animal populations. τ's contribution here is more modest but still meaningful: better contextualization of outbreak probability relative to environmental state.

8.4 Livestock heat stress, welfare, and productivity windows

Livestock heat-stress advisory systems currently offer three main types of outputs: temperature-humidity index (THI) thresholds for dairy cattle, coarse heat-alert categories for poultry and pigs, and general extreme-heat advisories tied to synoptic weather events. These outputs are useful as far as they go, but they do not provide species-specific, production-stage-specific, microclimate-aware, or operationally specific guidance.

Under τ, the operational gains would be: - Species-specific and production-stage-specific warnings (e.g., distinguishing risk windows for late-lactation dairy cows from dry cows, or for broilers in dense housing from free-range flocks). - Microclimate-aware risk windows that account for shelter, shade availability, water access, and local topography. - Transport and handling windows: specific go/no-go recommendations for livestock movement on days when heat exposure during transit could cause mortality or welfare compromise. - Expected production and welfare loss envelopes that allow farm managers and cooperatives to plan cooling, feed, and veterinary resource deployment in advance. - Integration with power and water contingency planning for cooling systems.

This is likely one of the fastest routes from τ capability to farmer-visible benefit in this entire portfolio, because the physical problem — predicting local heat exposure given a weather forecast — is well-defined and the operational response actions are clear.

8.5 One Health fusion: plant, animal, human, and environment

The deepest structural opportunity in this portfolio is not any single use case but the integration of plant-health, animal-health, and human-health biosecurity into a single causal weather–ecology substrate.

FAO's recent plant-health work explicitly frames plant health as a One Health issue: plant diseases affect not only food yields but human nutrition, biodiversity, ecosystem services, and the scale and nature of chemical interventions in agricultural landscapes. The environmental externalities of pesticide over-use ripple into animal health, water quality, soil biology, and human exposure. Managing plant health well is inseparable from managing the broader agro-ecological system.

A τ early-warning twin that supports plant-health, animal-health, and heat-stress advisory from the same physical substrate would provide a genuinely One Health-compatible platform: one in which a compound event — drought plus pest outbreak plus heat plus vector emergence — is processed through a single integrated risk engine rather than four separate institutional pipelines. This is the category of capability that no incumbent currently provides.

9 Geographic Case Studies

9.1 Case Study 1: Desert Locust — East Africa and Horn of Africa, 2020–2021

Context. The 2019–2021 desert locust outbreak was the most severe in 25 years for East Africa and 70 years for the Horn of Africa. At its peak, swarms were recorded in Djibouti, Ethiopia, Kenya, Somalia, Sudan, Yemen, India, and Pakistan. Individual swarms covered up to 2,400 square kilometers. FAO estimated in March 2020 that more than 25 million people in East Africa alone faced acute food-security threats directly linked to locust damage to crops and pasture.

Economic scale. FAO's emergency response in 2020 cost approximately USD 120 million. In the most severely affected areas — parts of Ethiopia, Somalia, and Kenya — crop and pasture losses in a single locust season exceeded 80 percent of seasonal production. Pastoralist communities in northern Kenya and southern Ethiopia lost primary grazing areas entirely; the secondary livestock-sector losses were not separately quantified but were substantial.

Baseline problem. FAO DLIS provided six-week locust outlooks throughout the outbreak and was the operational backbone of the international response. The system performed well at the broad regional level — identifying affected countries, tracking the general movement of swarms, and providing rainfall-based habitat-suitability maps. However, several operational gaps were documented in FAO's own post-event assessments:

- Swarm-location forecasting at the sub-national and sub-district level was insufficiently precise to target aerial treatment operations without large margins of uncertainty.
- Ground response was delayed in conflict-affected areas of southern Somalia and parts of Yemen, partly because imprecision in swarm-location forecasts made it difficult to prioritize limited access windows.
- Rainfall and vegetation models identified large potential breeding areas, but did not reliably distinguish zones of active breeding from zones of favorable conditions without breeding — leading to some unnecessary treatment deployments.
- Integration between locust risk and livestock-sector impact assessment was manual and lagged, delaying feed-emergency planning for pastoralist communities whose pasture was at risk.

τ-enabled change. Under the τ assumptions, the following operational improvements are plausible:

A law-faithful discrete twin of atmospheric transport, rainfall pulse dynamics, and vegetation response would improve habitat-suitability forecasting at the sub-district level by reducing the uncertainty envelope around rainfall-triggered breeding predictions. Swarm genesis probability maps would more reliably distinguish active breeding zones from merely favorable zones, allowing treatment aircraft to be directed with less wasted coverage. Outbreak-front tracking as swarms move across corridors would improve the lead time of advance warnings to districts not yet infested.

Integration with livestock-sector impact models — pasture availability, water access, grazing corridor conditions — would allow simultaneous generation of locust-risk and livestock-emergency advisories from the same weather substrate, replacing the sequential, manually integrated workflow used in 2020. This integration is not optional; in the Sahel and Horn of Africa, where crop farmers and pastoralists occupy overlapping landscapes, the decision about whether to trigger a feed-emergency response for livestock is inseparable from the decision about where to concentrate locust treatment.

Reference organizations. FAO DLIS (operational lead), IGAD (Intergovernmental Authority on Development, regional coordination), USAID (emergency response funding), ACMAD (African Centre of Meteorological Application for Development, weather services), national locust control organizations in Ethiopia, Kenya, Somalia, Sudan.

9.2 Case Study 2: Fall Armyworm — Nigeria and Ghana, 2016–Present

Context. Fall armyworm (Spodoptera frugiperda), native to the Americas, was first detected in Africa in 2016 in Nigeria and Ghana. By 2018 it had spread to 44 African countries. By the early 2020s it had established in South Asia and parts of Southeast Asia. FAO estimates potential losses of USD 13 billion per year across affected maize-growing areas in Africa alone. In sub-Saharan Africa, maize provides 60 to 70 percent of caloric intake; the welfare implications of inadequate early warning are severe.

Baseline problem. FAMEWS provides a real-time global and sub-national infestation map built on scouting and pheromone-trap data from national plant-protection organizations. The platform has been a valuable institutional achievement — it creates a shared data layer where none existed before — but its limitations are structural:

- FAMEWS is primarily a surveillance reporting platform, not a predictive forecast engine. It shows where fall armyworm has been observed, not where it is likely to emerge next based on weather dynamics and phenology.

- The link between national infestation maps and local farmer action is weak in many countries. Extension services are insufficiently resourced to translate national-level FAMEWS analytics into field-level scouting priorities and timely treatment recommendations.
- Chemical use in response to fall armyworm in West Africa has been intense, often untargeted, and in some areas counterproductive — generating resistance, harming beneficial insects, and increasing production costs without proportionate yield protection.
- Biological control interventions (e.g., Bacillus thuringiensis products) require application at the correct crop stage and weather window to be effective; the absence of precise emergence prediction reduces the efficacy of these lower-impact alternatives.

τ-enabled change. Under the τ assumptions, the operational improvements would be:

A physically faithful model of fall armyworm emergence, development, and dispersal would link temperature and rainfall sequences to generation timing, larval development windows, and adult migration distances — allowing district-level outbreak probability maps to be generated ahead of visible infestation, not after. Scouting recommendations would be targeted to specific fields and windows where risk is highest, reducing the total scouting burden while improving detection probability.

Treatment windows would be flagged with greater precision: the τ substrate would track the crop-stage and weather-window intersection that determines whether a biological control intervention is likely to be effective, reducing blanket chemical applications and enabling more rational use of lower-impact alternatives. Over a season, the combination of targeted scouting, better-timed biological control, and reduced blanket chemical applications could meaningfully reduce both yield losses and pesticide expenditure for smallholder farmers who currently bear both risks simultaneously.

Reference organizations. FAO FAMEWS (monitoring platform), CABI Plantwise Nigeria (extension and diagnostic support), USDA ARS (armyworm biology research), CGIAR-CIMMYT (integrated pest management research), Ghana Ministry of Agriculture (national response coordination).

9.3 Case Study 3 (Supplementary): Rift Valley Fever — Kenya, 2023–2024

Context. A Rift Valley Fever (RVF) outbreak in Kenya in late 2023 and early 2024 was triggered by above-normal rainfall in December 2023 and January 2024. Human infections and livestock casualties were reported across multiple counties in Kenya, with small numbers of human fatalities. RVF is a zoonosis transmitted by Aedes mosquitoes whose breeding is strongly correlated with heavy rainfall and flooding events in dryland landscapes.

Baseline problem. FAO operates a dedicated RVF decision-support tool that can issue alerts one to two months before first infections are reported in countries at risk, based on rainfall and vegetation condition indicators associated with Aedes breeding habitat. This capacity existed during the lead-up to the 2023 Kenya outbreak; post-event assessment indicated that the tool was not fully operationalized for real-time use in the weeks preceding the outbreak, and that the link between the environmental signal, vaccination campaign mobilization, and livestock movement restriction was not activated with sufficient lead time.

τ-enabled change. A τ early-warning twin would couple real-time rainfall and flooding data directly to a physical model of Aedes mosquito breeding habitat probability, providing continuous, spatially resolved risk maps throughout the rainy season rather than periodic alert products. Livestock vaccination campaign windows — the most effective preventive intervention — would be flagged weeks in advance, coordinated with the One Health channel to simultaneously alert human health authorities about elevated human-exposure risk in affected areas. The coordination layer between environmental signal, livestock health response, and human health response is precisely the domain where a shared physical substrate would have the greatest institutional value.

10 Finance, ROI, and Climate-Finance Eligibility

10.1 Benefit-cost reference from existing systems

FAO's 2020 desert locust response provides a useful benchmark for the economics of early-warning investment. FAO's emergency response cost approximately USD 120 million. FAO's own assessment estimated that timely control operations prevented losses of more than USD 500 million — an implicit benefit-cost ratio of approximately 4:1 for the response program as a whole. FAO's broader locust control program literature reports benefit-cost ratios ranging from 3:1 (late response, large outbreak already established) to 15:1 (early response, swarms controlled before reaching high density).

These ratios are conservative in one important respect: they count direct crop and pasture losses avoided but do not fully capture secondary benefits — avoided food price spikes, avoided livelihood shocks, avoided humanitarian response costs, and avoided long-term soil and vegetation recovery costs from massive locust damage. If those secondary effects are included, the economic case for investing in improved early warning is substantially stronger than even 15:1.

The analogous case for animal disease is less precisely documented at the system level but structurally similar. A USD 30 million investment in improved RVF early warning across East Africa that enabled earlier vaccination campaigns and reduced livestock mortality by even 1 percent of the annual regional burden would generate benefits substantially exceeding its cost.

10.2 Cost scenarios for τ integration

Scenario 1: National integration — single country, one to two priority pest or disease systems. A national-level integration of τ early warning into existing plant-health and animal-health surveillance systems, operating in shadow mode for one to two priority systems and generating advisory products for extension delivery, would cost approximately USD 5 to 12 million per country across a three-year Phase 1 period. This covers τ platform development, local data integration, shadow-mode benchmarking against FAO/WOAH systems, extension delivery infrastructure, and local capacity building. Countries with strong institutional anchors — an active FAO country program, a functional national plant-protection organization, a veterinary surveillance network — are the best first candidates.

Scenario 2: Regional corridor platform — East Africa/Sahel biosecurity platform. A regional platform serving 15 or more countries across the East Africa and Sahel corridor, integrating locust, fall armyworm, Rift Valley Fever, and livestock heat-stress advisories in a shared τ substrate, would cost approximately USD 30 to 60 million over a five-year Period. This would cover shared cloud infrastructure, regional data ingestion from FAO/WOAH/WMO systems, national extension integration points for each participating country, regional coordination mechanisms with FAO DLIS and FAMEWS, and ongoing model validation and calibration. The scale is consistent with comparable regional agricultural risk platforms and reflects the cost of building a genuinely operational, multi-country biosecurity layer.

10.3 Climate-finance eligibility

World Bank One Health Platform. The World Bank has committed approximately USD 2 billion across IDA and IBRD lending to One Health systems across multiple countries in the post-COVID period. The One Health Platform explicitly supports integrated animal, human, and environmental health system strengthening. A τ early-warning twin that delivers One Health-integrated biosecurity intelligence — plant, animal, vector, and human health channels from a shared physical substrate — is strongly aligned with this lending category.

Green Climate Fund — Anticipatory Action and Adaptation Window. GCF's Resilience for the Most Vulnerable category supports anticipatory action against climate-driven risks. A τ biosecurity platform that provides earlier, more targeted warnings for climate-sensitive threats — desert locust outbreaks triggered by rainfall anomalies, RVF outbreaks triggered by flooding events, livestock heat stress from temperature extremes — fits this category directly. The GCF has an established track record of funding early-warning system investments in agriculture, particularly in LDC and SIDS contexts.

FAO Emergency and Resilience Trust Fund (ERTF). The ERTF supports anticipatory action, early warning, and resilience-building programs across FAO's mandate. Integration of τ into existing ERTF-funded plant-health and animal-health programs — particularly in locust-affected zones and fall armyworm corridors — would be eligible for ERTF support.

WOAH Global Health Emergency Fund. WOAH operates a dedicated Global Health Emergency Fund for animal-disease preparedness and outbreak response. The fund supports early-detection capacity, preparedness planning, and emergency response for animal diseases of international concern. A τ platform with demonstrated early-detection performance for diseases on WOAH's list — RVF, FMD, HPAI — would be a natural candidate for WOAH fund support.

Parametric insurance and anticipatory financing. The shift toward parametric agricultural insurance and anticipatory humanitarian finance creates a commercial and development finance channel for τ early warning. Parametric products require precise, verifiable indices — exactly the type of certified output τ is designed to provide. A τ-derived outbreak probability index, structured as a parametric trigger, could support agricultural insurance products for smallholder farmers in locust-prone or RVF-prone corridors, creating a revenue stream that partially offsets platform operating costs while providing direct risk-management value to farmers.

11 Evidence and Translation Ladder

11.1 Phase 1: Shadow Mode Against Current Public Systems (0–18 months)

Goal. Benchmark τ forecast outputs against existing operational public-system products without modifying official operations. Build the empirical evidence base needed for Phase 2 deployment authorization.

Priority benchmarks. - DLIS locust outlooks and swarm-location alerts for one active desert locust season in the East Africa corridor. - FAMEWS hotspot and spread analytics for fall armyworm across one full West African maize season. - FAO EMPRES-i+ and RVF decision-support tool outputs for one RVF-sensitive season in East Africa. - WOAH WAHIS official notification streams for two to three diseases of highest regional concern. - National livestock heat advisories in one dairy or poultry region with available historical weather and production data.

Operating rule. Compare τ outputs against incumbent outputs. Compare both against realized field outcomes. Publish error envelopes, lead-time comparisons, false-positive and false-negative rates, and operational-relevance assessments. No τ-specific operational advisory products are issued during Phase 1.

Phase 1 KPIs. Lead-time improvement quantified in days per system. False-positive and false-negative rates documented for each benchmark system. Spatial precision of hotspot forecasts quantified against realized field observations. At least one peer-reviewed or independently reviewed technical report published.

11.2 Phase 2: Regional Operational Pilots (18–36 months)

Goal. Run τ-enhanced early warning operationally in selected high-need, high-visibility regions, with formal institutional partnerships and public scorecard reporting.

Pilot portfolio. At minimum one pilot from each of the following categories: - One locust-sensitive dryland corridor, partnered with DLIS and relevant national locust control organizations. - One maize/fall-armyworm region, partnered with FAMEWS and national plant-protection organization. - One vector-borne animal-disease region (e.g., RVF-sensitive East Africa), partnered with FAO and national veterinary authority. - One dairy or poultry heat-stress region, partnered with a farmer cooperative or national livestock board. - One mixed crop–livestock district to test integrated One Health advisory delivery across all three channels.

Phase 2 KPIs. Avoided crop loss in pilot regions (quantified through end-of-season assessment). Avoided blanket pesticide application (hectares and volume). Earlier outbreak detection (days relative to official first report). Vaccination/movement-control timing gain (days relative to current practice). Heat-stress windows accuracy (versus realized production and mortality outcomes). Smallholder and women-farmer reach. Speed from warning generation to farmer-level advisory delivery.

11.3 Phase 3: Integration with Public Action Systems (36–72 months)

Goal. Integrate τ early warning into routine public risk-management operations: phytosanitary response planning, veterinary vaccination and movement-control campaigns, anticipatory financing triggers, public-procurement contingency planning, and humanitarian early-action protocols.

Institutional integration targets. - Formal interoperability agreements with FAO DLIS, FAMEWS, and EMPRES-i+. - Integration with WOAH WAHIS reporting workflows. - National-level adoption by plant-protection and veterinary authorities in at minimum five countries. - Integration with at least one parametric insurance or anticipatory-financing product. - Integration with at least one national or regional humanitarian early-action protocol.

Phase 3 KPIs. Number of countries with τ-integrated national advisory systems. Number of livelihoods reached by τ-informed advisories. Reduction in emergency intervention expenditure per country. Verified avoided losses attributed to earlier warnings. Trust and comprehension scores from farmer and extension surveys.

12 Stakeholder Map and Change Management

12.1 Primary institutional stakeholders

National plant-protection organizations (NPPOs). NPPOs are the national counterparts to FAO IPPC and carry statutory responsibility for phytosanitary surveillance and response. They are the natural operational entry point for τ integration on the plant-health side. Change management challenge: NPPOs have established workflows built around existing advisory products; τ integration must enhance rather than displace these workflows, and must be compatible with IPPC reporting obligations.

National veterinary authorities (NVAs). NVAs carry statutory responsibility for animal disease surveillance, notification, and response under WOAH obligations. They are the corresponding entry point for τ animal-health integration. Change management challenge: NVAs are often under-resourced and operate under significant reporting burden; τ must deliver decision-simplification, not additional reporting complexity.

FAO country and regional offices. FAO serves as the institutional anchor for both the DLIS/-

FAMEWS and EMPRES-i+ systems and has country presence in virtually all target geographies. FAO's alignment is essential for Phase 2 operational pilots and Phase 3 institutional integration.

WOAH regional representations. WOAH's five regional representations (Africa, Americas, Asia-Pacific, Europe, Middle East) are the coordination layer for national veterinary authority engagement. WOAH regional offices can facilitate simultaneous multi-country adoption of τ animal-health products.

12.2 Secondary and enabling stakeholders

National meteorological and hydrological services (NMHSs). NMHSs provide the weather data infrastructure on which τ operates. Early integration of NMHSs into pilot design — ensuring that τ can access real-time national weather observation networks — is critical to Phase 1 and Phase 2 performance.

Extension services and farmer organizations. Extension services are the last-mile delivery mechanism for τ advisories. Farmer cooperatives, commodity boards, and smallholder organizations are the end-user institutions. Engagement in pilot design is essential to ensure that τ outputs are translated into formats and delivery channels that farmers can act on.

Development finance institutions and bilateral donors. World Bank One Health Platform, GCF, USAID, EU development programs, and bilateral donors from Germany (BMZ), UK (FCDO), and France (AFD) are the primary funders of large-scale agricultural early-warning programs. Early engagement with these stakeholders in Phase 1 — sharing Phase 1 benchmark results — builds the funding case for Phase 2 and 3 deployment.

13 Gender, Equity, and Labor Dimensions

13.1 Women farmers and smallholder plant health

In sub-Saharan Africa, women constitute 60 to 80 percent of smallholder farmers, yet they are systematically underrepresented in plant-health surveillance systems, extension service contact lists, and national early-warning notification channels. When a fall armyworm or locust early warning is issued, the probability that it reaches a woman farmer in a remote district through official channels is substantially lower than the probability that it reaches a male farmer who is a registered cooperative member or extension contact.

τ-enabled advisory systems must explicitly address this gap. Advisory delivery pathways must reach women farmers directly — through mobile platforms in local languages, community radio, women's farmer group networks, and cooperative-linked messaging — rather than filtering through institutional channels that systematically under-represent women. Phase 1 pilot design should include disaggregated reach data by gender from the outset.

13.2 Occupational heat safety for agricultural labor

Agricultural workers are the most heat-exposed occupational group in most low- and middle-income countries. WMO reports that agricultural workers are 35 times more likely to die from occupational heat exposure than workers in other sectors. This is not simply a welfare statistic; it is a structural risk that falls disproportionately on seasonal, informal, and migrant agricultural labor — populations with the least institutional protection, the least access to cooling infrastructure, and the least ability to choose not to work in dangerous conditions.

τ-enabled livestock heat-stress and crop-work advisory systems should include explicit labor-safety dimensions: go/no-go windows for field operations (harvesting, irrigation, intensive animal husbandry), recommended rest-and-water schedules, and warnings calibrated to exposed outdoor worker populations rather than only to crop physiology or livestock welfare. These outputs should be integrated with national occupational health authorities and humanitarian labor-safety programs where relevant.

13.3 Smallholder livestock keepers and pastoral communities

Pastoral and agropastoral communities in dryland Africa and Central Asia are among the most exposed to compound biological hazards — locust outbreaks, RVF, and heat stress can occur simultaneously in the same landscape — and among the least served by existing advisory systems. Pastoral communities often lack addresses in formal notification systems, operate in areas with limited mobile network coverage, and have customary land management practices that may not be reflected in standard early-warning geographic units.

τ deployment in pastoral contexts requires deliberate adaptation: advisory products calibrated to pastoral decision cycles (movement timing, watering point planning, emergency destocking), delivery channels that reach mobile populations, and governance arrangements that respect community data rights and do not create surveillance infrastructure that could be used to restrict pastoral mobility without their consent.

14 Benchmark Suite and Success Metrics

Technical forecasting metrics. - Lead time for actionable warnings: target 7+ days advance of outbreak threshold crossing for migratory pests; 4+ weeks for vector-borne disease triggers; 3+ days for heat-stress events. - False-positive rate: fraction of warnings that do not correspond to realized events, by system and geography. - False-negative rate: fraction of realized events not preceded by a τ warning within the lead-time window. - Spatial precision: area of warning polygons versus area of realized infestation or risk, at district and sub-district scale. - Confidence envelope calibration: fraction of outcomes falling within the stated uncertainty range. - Performance under compound hazards: accuracy under simultaneous heat-plus-pest or drought-plus-disease events.

Plant-health metrics. - Avoided crop loss in pilot regions, quantified through end-of-season field assessment (tons per hectare and total value). - Reduction in blanket pesticide application events and volume per hectare in pilot areas. - Timeliness and specificity of scouting recommendations relative to realized outbreak timing and location. - Reduction in emergency response area per outbreak event.

Animal-health metrics. - Days of earlier outbreak detection relative to official WOAH notification date. - Vaccination and movement-control timing gain relative to current practice. - Diagnostic resource prioritization efficiency: fraction of laboratory capacity directed to highest-risk areas. - Avoided livestock mortality and morbidity in pilot systems, quantified through cooperative or veterinary service records.

Livestock heat-stress metrics. - Accuracy of heat-stress window predictions versus realized THI measurements at farm level. - Reduction in heat-linked livestock mortality events during pilot seasons. - Reduction in heat-linked production losses (milk, eggs, weight gain) in monitored pilot herds. - Reduction in heat-linked transport incidents and mortality during advisory-covered movement periods.

Equity and public-good metrics. - Total livelihoods reached by τ-informed advisories, disaggregated by gender and farm type. - Smallholder and pastoralist participation rate in pilot programs. - Reduction in loss burden per household in pilot areas relative to matched control areas. - Reduction in unnecessary chemical and emergency interventions. - Speed from warning generation to farmer-level advisory delivery (hours).

15 Governance Guardrails

Sound deployment of τ early warning in the plant-health and animal-health domain requires active governance attention to a set of structural risks. The following guardrails are essential.

1. Do not substitute overconfidence for field surveillance. A stronger physics engine should sharpen field surveillance targeting, not substitute for it. Ground truth — scouting reports, pheromone-trap data, field observations from farmers and extension workers — remains essential for model updating and outbreak confirmation. Any τ deployment that reduces field surveillance investment in the name of model confidence should be treated as a governance failure.

2. Avoid the "better warning → more blanket intervention" trap. The purpose of improved early warning is more targeted action, not simply a lower threshold for broad interventions. Governance frameworks must specify that improved warning precision is used to reduce unnecessary interventions, not to justify more frequent ones. This is particularly important for pesticide applications and movement restrictions, where overuse carries direct costs.

3. Maintain farmer and community agency in the advisory loop. Plant-health and animal-health advisory systems that centralize all interpretive authority will underperform socially even if they perform technically. Governance must protect the farmer's role as the primary decision-maker on their land and provide advisories as decision support, not directives.

4. Protect data rights and prevent punitive surveillance drift. Plant-health and animal-health data have direct implications for market access, trade eligibility, movement restrictions, and livelihood security. A τ early-warning platform that accumulates field-level disease or pest data must have robust data governance: clear rules about who owns the data, who can access it, for what purposes it can be used, and what protections exist against its use for punitive or restrictive purposes that harm the farmers and communities it was designed to serve.

5. Preserve One Health balance across channels. Because plant, animal, and environmental health are structurally intertwined, governance must prevent optimization of one channel in ways that damage the others. Reducing plant losses through excessive chemical use imposes animal-health, human-health, and environmental costs. Reducing animal-disease risk through aggressive movement restrictions imposes livelihood costs on pastoral and agropastoral communities. Governance frameworks must require cross-channel impact assessment before major advisory actions are taken.

6. Maintain chemical stewardship standards. Any τ-informed pesticide advisory must be embedded within integrated pest management principles: chemical intervention should be the last resort after biological, cultural, and physical control options have been evaluated. τ-improved precision creates an opportunity to reduce, not increase, overall pesticide use; governance must ensure that opportunity is realized rather than squandered.

7. Open validation before institutional lock-in. Given the stakes for food security, trade, and livelihoods, the correct deployment sequence is open benchmarking and transparent pilot evidence before any single-system institutional lock-in. Governments and international organizations should not be asked to replace existing advisory systems with τ-based alternatives before independent validation evidence is published. Phase 1 shadow-mode benchmarking serves this function explicitly.

8. Equity and inclusion audits at each phase transition. Transition from Phase 1 to Phase 2 and from Phase 2 to Phase 3 should be conditioned on evidence that τ advisory products are reaching underserved populations — women farmers, smallholder pastoralists, remote communities — at rates comparable to the general agricultural population. Phase transitions that cannot demonstrate equitable reach should be paused pending delivery infrastructure improvements.

16 SDG Mapping and Bottom Line

16.1 SDG alignment

SDG 2 (Zero Hunger). Plant-pest and crop-disease losses directly undermine food production; improved early warning reduces crop losses and supports food security for the most vulnerable populations. The 40 percent crop loss baseline and the USD 220+ billion trade-loss figure are direct SDG 2 metrics.

SDG 3 (Good Health and Well-being). Zoonotic and vector-borne diseases (RVF, bluetongue, HPAI) at the animal-human interface directly threaten human health. Occupational heat exposure in agriculture is a direct SDG 3 labor-safety issue. One Health integration across plant, animal, and human health channels connects τ to multiple SDG 3 targets simultaneously.

SDG 13 (Climate Action). Climate change is the primary driver of expanding pest ranges, shifting disease vectors, and increasing heat-stress frequency. τ early warning is explicitly a climate-adaptation tool: it reduces the economic and food-security costs of climate-driven biological hazards, making agricultural systems more resilient to a changing climate. GCF eligibility follows from this alignment directly.

SDG 15 (Life on Land). Reduced pesticide use through better-targeted early warning directly benefits terrestrial biodiversity: fewer non-target organism impacts, reduced chemical soil burden, and lower pesticide runoff into adjacent ecosystems. The explicit connection between plant-health early warning and pesticide-use reduction is FAO's own framing; τ extends that logic by making targeting more precise.

16.2 Bottom line

The burden addressed by this paper is large, well-documented, and growing. The institutional infrastructure for early warning already exists and is operational. The gap is not institution-building from scratch — it is improving the physical substrate on which existing institutions rely, and integrating currently separate plant-health, animal-health, and heat-stress channels into a unified causal engine.

If the τ assumptions hold well enough to matter operationally, this paper identifies a path to meaningful reduction of a combined burden that official institutions quantify in the range of USD 278–560 billion per year, across threats that directly affect the food security and livelihoods of billions of people. The finance ecosystem, the institutional partners, and the geographic pilots are ready. The question is whether the physics is sound enough to deliver. That is the question the Phase 1 shadow-mode benchmarks are designed to answer.

17 References

1. FAO (2025). *The hidden health crisis: How plant diseases threaten global food security.* FAO One Health Highlights. `https://www.fao.org/one-health/highlights/how-plant-diseases-threaten-global-food-security`
2. FAO (2025). *Transboundary animal diseases pose urgent threat to global food security, FAO warns.* FAO Newsroom, November 2025. `https://www.fao.org/newsroom/detail/transboundary-animal-diseases-pose-urgent-threat-to-global-food-security--fao-warns/en`
3. FAO/WMO (2025). *Extreme Heat and Agriculture: A joint FAO/WMO report on extreme heat risks for agrifood systems.* Rome and Geneva. `https://openknowledge.fao.org/server/api/core/bitstreams/f6506a63-fb04-4f34-bea6-cb92f665e94f/content`

4. WMO (2025). *FAO and WMO report highlights extreme heat risks for agriculture.* WMO News. https://wmo.int/media/news/fao-and-wmo-report-highlights-extreme-heat-risks-agriculture
5. FAO DLIS. *Desert Locust Information Service — Locust Watch.* 24/7 monitoring, six-week outlooks, and six-month precipitation predictions. https://www.fao.org/locust-watch/activities/dlis-home/dlis/en
6. FAO Locust Watch. *Desert Locust Information Service — main portal.* https://www.fao.org/locusts/en/
7. FAO FAMEWS. *Fall Armyworm Monitoring and Early Warning System — FAMEWS Global Platform.* Real-time infestation maps and analytics. https://www.fao.org/fall-armyworm/monitoring-tools/famews-global-platform/en
8. FAO. *Fall armyworm monitoring tools overview.* https://www.fao.org/fall-armyworm/monitoring-tools/en/
9. FAO. *Early Warning, Early Action through FAO EMPRES-i+.* Animal health area of work. https://www.fao.org/animal-health/areas-of-work/early-warning-and-disease-intelligence/Early-Warning-Early-Action-through-FAO-EMPRES-i-/en
10. FAO. *Early warning and disease intelligence — animal health overview.* https://www.fao.org/animal-health/areas-of-work/early-warning-and-disease-intelligence/en
11. WOAH. *World Animal Health Information System (WAHIS).* Early warning, monitoring, and official notification platform. https://www.woah.org/en/what-we-do/animal-health-and-welfare/disease-data-collection/world-animal-health-information-system/
12. WOAH (2025). *The State of the World's Animal Health 2025.* World Organisation for Animal Health, Paris. https://www.woah.org/en/the-state-of-the-worlds-animal-health/
13. FAO (2025). *How early warning systems for plant health protect crops, people, and planet.* FAO One Health Highlights. https://www.fao.org/one-health/highlights/early-warning-systems-for-plant-health/en
14. FAO (2020). *Desert Locust crisis in East Africa — emergency response updates and situation reports.* FAO Locust Watch. https://www.fao.org/locusts/en/
15. FAO (2018). *Fall armyworm spread and impact in sub-Saharan Africa — assessment of the situation and actions required.* FAO, Rome.
16. CABI. *Plantwise programme and plant clinic network.* Plant health advisory and digital diagnostics across ~90 countries. https://www.cabi.org/projects/plantwise/
17. Cropin. *Smart farming and disease intelligence platform.* Crop monitoring, disease forecasting, satellite and weather integration. https://www.cropin.com/
18. PEAT GmbH. *Plantix — mobile plant disease diagnostic platform.* AI image-based disease identification, 10M+ downloads. https://plantix.net/
19. SatAgro. *Satellite-based crop monitoring and precision agriculture platform.* NDVI and weather integration, Eastern Europe focus. https://satagro.pl/
20. WOAH. *Global Health Emergency Fund for animal diseases and preparedness.* World Organisation for Animal Health, Paris. https://www.woah.org/
21. World Bank. *One Health Platform and IDA/IBRD lending for integrated animal-human-environmental health systems.* Approximately USD 2 billion committed post-COVID. https://www.worldbank.org/en/topic/agriculture/brief/one-health
22. Green Climate Fund (GCF). *Anticipatory action and adaptation priorities — Resilience for the Most Vulnerable category.* https://www.greenclimate.fund/

23. FAO. *Emergency and Resilience Trust Fund (ERTF)*. FAO anticipatory action, early warning, and resilience-building programs. `https://www.fao.org/emergencies/resources/documents/resources-detail/en/c/1247816/`

24. IGAD. *Intergovernmental Authority on Development — locust response coordination, Horn of Africa*. Djibouti. `https://igad.int/`

25. USAID. *Africa locust emergency response 2020 — USAID emergency contributions and operational support*. Washington, DC.

26. CGIAR-CIMMYT. *Fall armyworm integrated management research across Africa and South Asia*. `https://www.cimmyt.org/`

27. FAO/CABI. *Spodoptera frugiperda: fall armyworm integrated pest management assessment*. FAO-CABI joint assessment of biological control, chemical use, and management options.

28. FAO (2022). *Desert Locust 2019–2021 crisis — post-event operational review and lessons learned*. FAO Locust Watch, Rome.

29. ACMAD. *African Centre of Meteorological Application for Development — weather services for locust and agricultural risk in Africa*. Niamey, Niger. `https://www.acmad.net/`

30. FAO GLEWS+. *Global Early Warning and Response System for Major Animal Diseases including Zoonoses*. `https://www.fao.org/animal-health/areas-of-work/early-warning-and-disease-intelligence/en`

31. FAO. *Rift Valley Fever decision support tool — early warning for high-risk areas and alert timelines*. In: FAO Animal Health early warning documentation.

32. WOAH. *Rift Valley Fever disease card and monitoring data*. WAHIS listed disease surveillance. `https://www.woah.org/en/disease/rift-valley-fever/`

33. WMO (2022). *State of Climate Services: Agriculture and Food Security 2022*. Geneva. `https://public.wmo.int/en/resources/library/2022-state-of-climate-services`

34. FAO (2023). *Agricultural workers' exposure to pesticides: occupational health and safety assessment*. FAO, Rome.

35. IPCC (2022). *Climate Change 2022: Impacts, Adaptation and Vulnerability. Chapter 5: Food, Fibre, and Other Ecosystem Products*. IPCC Sixth Assessment Report, Working Group II.

36. Bebber, D.P., Ramotowski, M.A.T., Gurr, S.J. (2013). Crop pests and pathogens move polewards in a warming world. *Nature Climate Change*, 3, 985–988.

37. Deutsch, C.A. et al. (2018). Increase in crop losses to insect pests in a warming climate. *Science*, 361(6405), 916–919.

38. FAO (2021). *The impact of disasters and crises on agriculture and food security 2021*. FAO, Rome. `https://www.fao.org/3/cb3673en/cb3673en.pdf`

39. World Bank (2024). *One Health Investment Framework: Strengthening integrated animal, human, and environmental health systems in low- and middle-income countries*. World Bank Group, Washington, DC.

40. WOAH (2024). *Terrestrial Animal Health Code, Volume I: General Provisions, Chapter 3.4 — Early warning and emergency management*. 32nd Edition. Paris.

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 weather-pathogen-vector-host twin could sharpen plant-health, animal-disease, and livestock heat-stress early warning across the full agricultural biosecurity stack. The v3 impact thesis is conditional: a Tau-grade weather-pest-disease-livestock stress early-warning 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 weather-pathogen-vector-host twin could sharpen plant-health, animal-disease, and livestock heat-stress early warning across the full agricultural biosecurity stack. 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 weather-pest-disease-livestock stress early-warning 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: **weather-pest-disease-livestock stress early-warning twin**.

First benchmarkable test: outbreak, vector, heat-stress, and livestock-risk alerts against public surveillance and advisory 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 Pest, Disease, and Livestock-Stress Early Warning without operational authority.	medium
Realistic	A reviewed prototype strengthens several public-sector workflows for Pest, Disease, and Livestock-Stress Early Warning after benchmark comparison with incumbent systems.	medium-low
Optimistic	A reusable public-good intelligence layer becomes plausible for Pest, Disease, and Livestock-Stress Early Warning 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

district early-warning pilot linking meteorological, veterinary, crop-protection, and extension agencies

18.10 Benchmark suite and success metrics

Type	Incumbent baseline	Required Tau benchmark	Success metric	Validator
translation benchmark	current public or institutional systems in the domain	outbreak, vector, heat-stress, and livestock-risk alerts against public surveillance and advisory 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 pest, disease, and livestock-stress early warning. `https://panta-rhei.site/impact/papers/pest-disease-livestock-stress-early-warning/`, 2026. Current public full-text source for dossier pest-disease-livestock-stress-early-warning.

[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/pest-disease-livestock-stress-early-warning/`, 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 Pest, Disease, and Livestock-Stress Early Warning

Dossier ID: PGID-AGRI-03 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