Panta Rhei
Research Program

Research Briefings · Public-Good Impact Dossiers

One Health · Food, Life & Health Systems

τ for Precision Public Health, Physiology-Aware Prevention, and Safer Therapeutics / Response Optimization

Conditional public-good pathway for Precision Public Health, Physiology-Aware Prevention, and Safer Therapeutics / Response Optimization

Public-Good Impact Dossier

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

Dossier ID: PGID-OH-04 Release: May 2026 publication-ready release

Thorsten Fuchs • Anna-Sophie Fuchs

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

Safe harbor / release discipline

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

Briefing identity and source routes

Dossier ID	PGID-OH-04
Portfolio	One Health
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/precision-public-health-physiology-prevention-safer-therapeutics/
Landing route	https://panta-rhei.site/publications/research-briefings/public-good/precision-public-health-physiology-prevention-safer-therapeutics/

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

Preventable harm at population scale remains one of the most tractable remaining frontiers in global health. WHO data confirm that **noncommunicable diseases killed at least 43 million people in 2021**, including **18 million premature deaths before age 70**, with 82% of those premature deaths occurring in low- and middle-income countries.[1] Medication-related harm, which accounts for **nearly half of all preventable harm in medical care**, costs the global health system an estimated **US$42 billion per year** in avoidable adverse drug events, with the toll falling disproportionately on older adults, polypharmacy patients, and populations with inadequate monitoring infrastructure.[2][3][4] Climate-sensitive health emergencies — extreme heat, air pollution episodes, vector-borne disease surges, and infectious outbreaks — kill hundreds of thousands of additional people each year under predictive conditions that current surveillance systems cannot adequately anticipate.[5][6]

The institutional will to address this exists. In February 2026, the WHO Executive Board formally advanced precision medicine as a global governance topic, recommending a World Health Assembly resolution on **precision medicine as a path towards targeted, personalized, and equitable care**, explicitly tying it to prevention, diagnosis, and treatment across the life course and to predictive, preventive, and participatory health systems.[7] WHO Africa's **Precision Public Health Strategy 2024–2030** has already set regional targets for data-driven disease control and tailored interventions.[8] ARPA-H's **CATALYST** programme is funding human-based computational physiology models to improve drug safety prediction.[9] The US FDA continues to expand its **model-informed drug development (MIDD)** regulatory frameworks to reduce therapeutic failure rates.[10][11]

What remains missing is a stronger causal engine: a physically grounded, bounded-error, multi-scale twin of human physiology that can couple environmental exposures, body state dynamics, treatment interactions, and population heterogeneity into a single coherent predictive substrate. That is precisely the opening this paper addresses.

Under the working assumption that the τ (Category τ) framework delivers a **law-faithful, bounded-error, coarse-grainable discrete twin** of physiological dynamics, environmental-exposure coupling, and health-system state evolution, this paper argues that the following consequences follow for precision public health, prevention, physiology-aware care, and response optimization:

- Heat-health surveillance extends from 3-day reactive alerts to 7–10-day anticipatory planning

[1] WHO. *Noncommunicable diseases — fact sheet.* Updated April 2025. https://www.who.int/news-room/fact-sheets/detail/noncommunicable-diseases

[2] GBD 2019 Diseases and Injuries Collaborators. *Global burden of 369 diseases and injuries in 204 countries and territories, 1990–2019: a systematic analysis for the Global Burden of Disease Study 2019.* Lancet. 2020;396(10258):1204–1222. DOI: 10.1016/S0140-6736(20)30925-9

[3] WHO. *Medication without harm: Policy brief.* 2024. https://www.who.int/publications/i/item/9789240062764

[4] WHO. *The third global patient safety challenge: tackling medication-related harm.* Geneva, 2017. https://iris.who.int/bitstream/handle/10665/255263/WHO-HIS-SDS-2017.6-eng.pdf

[5] Robine JM, Cheung SL, Le Roy S, et al. *Death toll exceeded 70,000 in Europe during the summer of 2003.* C R Biol. 2008;331(2):171–178. DOI: 10.1016/j.crvi.2007.12.001

[6] Ballester J, Quijal-Zamorano M, Méndez Turrubiates RF, et al. *Heat-related mortality in Europe during the summer of 2022.* Nat Med. 2023;29:1857–1866. DOI: 10.1038/s41591-023-02419-z

[7] WHO Executive Board. *Precision medicine: a path towards targeted, personalized and equitable care (EB158.R2).* 2026. https://apps.who.int/gb/ebwha/pdf_files/EB158/B158_R2-en.pdf

[8] WHO Regional Office for Africa. *Precision public health strategy, 2024–2030.* 2024. https://www.afro.who.int/publications/precision-public-health-strategy-2024-2030

[9] ARPA-H. *CATALYST — Computational ADME-Tox and Physiology Analysis for Safer Therapeutics.* 2024. https://arpa-h.gov/explore-funding/programs/catalyst

[10] FDA. *Model-Informed Drug Development Paired Meeting Program.* Updated 2026. https://www.fda.gov/drugs/development-resources/model-informed-drug-development-paired-meeting-program

[11] FDA. *Focus Area: Model-Informed Product Development.* https://www.fda.gov/science-research/focus-areas-regulatory-science-report/focus-area-model-informed-product-development

windows, enabling pre-positioning of elder care and hospital surge capacity ahead of events.
- Malaria and dengue climate-responsive early warning systems gain 8–12 weeks of predictive lead time, enabling anticipatory drug and net pre-positioning before transmission windows open.
- Medication-safety systems acquire the physiological-environmental coupling needed to flag heat-sensitive polypharmacy risk, transition-of-care failure points, and population-heterogeneous dose-response profiles.
- Hospital surge prediction for infectious outbreaks shifts from 50–80% uncertainty at two weeks to 15–25% uncertainty, enabling decisive ventilator and staffing decisions.
- Population-level NCD prevention targeting becomes risk-stratified rather than demographic-average, improving screening efficiency and early-intervention uptake in high-burden subgroups.

The paper is structured for **ministries of health, public-health institutes, hospital systems, medicines regulators, health-technology assessment bodies, philanthropic health funders, and digital-health governance teams**. It does not claim that the broader scientific community has validated the τ framework for these applications. It asks what would follow **if** those capabilities were operational, and what governance, investment, and deployment design would be required to realize them responsibly.

2 Why This Matters Now

Four converging pressures make 2026 a strategically important moment for deploying improved decision-intelligence tools in public health and clinical systems.

The preventable-harm burden is enormous, and the tools to address it are not yet adequate. WHO's 2026 precision medicine resolution does not frame precision approaches as a niche offering for high-income settings. It frames them as an important enabler of **targeted, effective, and efficient interventions** that can strengthen universal health coverage and help ensure that the right intervention reaches the right patient at the right time.[12] That institutional reorientation reflects a real gap: most prevention and treatment systems still operate with population averages that do not capture the physiological and environmental heterogeneity that drives differential outcomes. The consequence is familiar — missed prevention windows, avoidable adverse drug events, costly trial-and-error treatment selection, and widening inequity whenever high-resolution medicine is accessible only in well-resourced settings.

Climate change is converting many health risks from chronic-slow to acute-episodic, compressing response windows. The 2003 European heat wave killed more than 70,000 people across the continent in a single summer.[13] The 2022 heat event killed an estimated 61,672 people across Europe, even after two decades of adaptation.[14] Malaria transmission windows in highland East Africa are lengthening and shifting seasonally in ways that 4–6 week detection lags cannot capture.[15] Air quality-respiratory health relationships are being reshaped by the convergence of urban pollution, wildfire smoke, and ambient ozone in ways that current epidemiological early warning systems cannot model in real time.[16] A health intelligence infrastructure that cannot match these dynamics with adequate predictive lead time is structurally vulnerable to the climate-sensitive health burden of the coming decades.

[12]WHO Executive Board. *Precision medicine: a path towards targeted, personalized and equitable care (EB158.R2).* 2026. https://apps.who.int/gb/ebwha/pdf_files/EB158/B158_R2-en.pdf

[13]Robine JM, Cheung SL, Le Roy S, et al. *Death toll exceeded 70,000 in Europe during the summer of 2003.* C R Biol. 2008;331(2):171–178. DOI: 10.1016/j.crvi.2007.12.001

[14]Ballester J, Quijal-Zamorano M, Méndez Turrubiates RF, et al. *Heat-related mortality in Europe during the summer of 2022.* Nat Med. 2023;29:1857–1866. DOI: 10.1038/s41591-023-02419-z

[15]WHO. *World Malaria Report 2022.* Geneva: World Health Organization, 2022. https://www.who.int/teams/global-malaria-programme/reports/world-malaria-report-2022

[16]WHO. *Ambient (outdoor) air pollution — fact sheet.* 2023. https://www.who.int/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health

Official institutions are already moving toward computational physiology — but the physical grounding is still missing. ARPA-H's CATALYST programme is explicitly trying to build **human-based computational models of physiology** that can predict drug safety and efficacy in silico.[17] FDA's model-informed product development work signals that regulators are ready to accept computational evidence where it meets quality thresholds.[18][19] WHO Africa's precision public health strategy calls explicitly for advanced analytics, geoscience, and mathematical modelling to improve disease burden estimates and identify susceptible sub-national units.[20] The institutional pathway already exists. The missing element is a physically faithful causal substrate that couples environmental, physiological, and intervention dynamics in a single coherent engine.

The equity imperative has become a governance baseline, not a secondary consideration. WHO's 2026 resolution is explicit about underrepresentation in biomedical datasets, the digital divide, inequitable access to precision care, and the risk that precision medicine deepens rather than narrows global health disparities.[21] Any deployment that does not address this dimension head-on will face legitimate institutional resistance and will fail its most important public-good test. A τ-grade approach that cannot be deployed for under-resourced populations in sub-Saharan Africa, South Asia, and Oceania is not a public-interest tool — it is a premium-market offering wearing a development-sector label.

Together, these pressures define a window: the clinical and public-health communities already know the direction they need to travel, the institutional frameworks already exist, the political will is building — what is missing is the physics-faithful, multi-scale, environmentally coupled causal engine that makes the mission executable rather than aspirational.

3 Scope and Reader Orientation

This is **Paper 4 of 4** in the Panta Rhei Impact One Health companion series, and specifically the **precision public health and therapeutic safety** tier of that portfolio. It covers:

- precision public health and targeted prevention through risk-stratified subpopulation intelligence;
- physiology-aware prevention and care pathways including heat, air-quality, and environmental-exposure coupling;
- medication safety and therapeutic-response modeling under population and environmental heterogeneity;
- model-informed safer therapeutics and drug-development acceleration;
- hospital surge prediction and climate-responsive health-system response optimization;
- and equity-aware, public-interest implementation across high- and lower-resource settings.

This paper does not revisit the topics covered in earlier One Health papers. Paper 1 covers vector-borne disease, zoonotic spillover, and climate-sensitive outbreak early warning. Paper 2 covers health-system resilience, facility continuity, cold chains, and clinical operations infrastructure. Paper 3 covers AMR, wastewater/environmental surveillance, and environmental transmission intelligence. This paper — Paper 4 — adds the human-physiology and therapeutic layer: the layer that translates environmental and epidemiological intelligence into precise, individualized, and population-stratified

[17] ARPA-H. *CATALYST — Computational ADME-Tox and Physiology Analysis for Safer Therapeutics.* 2024. https://arpa-h.gov/explore-funding/programs/catalyst

[18] FDA. *Model-Informed Drug Development Paired Meeting Program.* Updated 2026. https://www.fda.gov/drugs/development-resources/model-informed-drug-development-paired-meeting-program

[19] FDA. *Focus Area: Model-Informed Product Development.* https://www.fda.gov/science-research/focus-areas-regulatory-science-report/focus-area-model-informed-product-development

[20] WHO Regional Office for Africa. *Precision public health strategy, 2024–2030.* 2024. https://www.afro.who.int/publications/precision-public-health-strategy-2024-2030

[21] WHO Executive Board. *Precision medicine: a path towards targeted, personalized and equitable care (EB158.R2).* 2026. https://apps.who.int/gb/ebwha/pdf_files/EB158/B158_R2-en.pdf

prevention and care decisions.

Working assumption. All claims about τ-enabled improvements throughout this paper are conditioned on the assumption that the τ framework delivers law-faithful, bounded-error, coarse-grainable twins of physiological dynamics, environmental-exposure coupling, and health-system state evolution. This is a planning paper, not a validation report. Readers should evaluate the argument conditional on that assumption.

Caveats are structural. This paper separates clearly between: - what official institutions already know and want from health intelligence systems; - what τ would newly provide under the working assumption; - and what impact scenarios are reasoned planning inferences rather than official forecasts or validated experimental results.

The target reader is not being asked to accept τ's mathematical claims. They are being asked to evaluate: **if this capability were real, would it change what you do, for whom, and at what scale?**

4 The Opportunity Baseline

The public-good case for precision public health and safer therapeutics is already established in official data. Five quantitative anchors define the baseline against which τ-grade improvements should be measured.

Noncommunicable disease burden: 43 million deaths per year, 18 million premature. WHO's current NCD fact sheet attributes at least 43 million deaths per year to noncommunicable diseases, including cardiovascular disease (17.9 million), cancer (9.3 million), chronic respiratory disease (4.1 million), and diabetes (2 million).[22] Of these, 18 million occur before age 70 — premature deaths that are, in most cases, partially preventable through better risk identification, earlier intervention, and more physiologically calibrated treatment. The concentration of premature NCD deaths in LMICs (82% of the total) underlines that the marginal value of better prevention targeting is highest where health systems are least equipped to compensate for poor targeting with high-cost reactive care.

Medication harm: US$42 billion per year, nearly half of all preventable harm. WHO's Medication Without Harm initiative estimates annual global costs from medication errors at approximately US$42 billion.[23] A 2024 WHO policy brief states that harm due to medicines and therapeutic options accounts for **nearly 50% of preventable harm in medical care**.[24] This is not primarily a prescribing-error problem in the narrow sense: it is a physiological heterogeneity problem. The same drug regimen that is safe and effective for a 45-year-old with normal renal function can be harmful for a 75-year-old with declining kidney function under summer heat. The same polypharmacy regime that is manageable in winter becomes dangerous when thermoregulatory stress or dehydration shifts drug clearance rates. Current systems are not designed to track these dynamics continuously across large populations.

Climate-sensitive mortality: 70,000+ from heat in Europe, 619,000 from malaria globally. The 2003 European heat wave caused more than 70,000 excess deaths (Robine et al. 2008).[25] The 2022 heat wave caused 61,672 excess deaths across Europe (Ballester et al., Nature Medicine, 2023), demonstrating that even after two decades of adaptation and alert system investment, the mortality

[22] WHO. *Noncommunicable diseases — fact sheet.* Updated April 2025. https://www.who.int/news-room/fact-sheets/detail/noncommunicable-diseases

[23] WHO. *The third global patient safety challenge: tackling medication-related harm.* Geneva, 2017. https://iris.who.int/bitstream/handle/10665/255263/WHO-HIS-SDS-2017.6-eng.pdf

[24] WHO. *Medication without harm: Policy brief.* 2024. https://www.who.int/publications/i/item/9789240062764

[25] Robine JM, Cheung SL, Le Roy S, et al. *Death toll exceeded 70,000 in Europe during the summer of 2003.* C R Biol. 2008;331(2):171–178. DOI: 10.1016/j.crvi.2007.12.001

burden from heat events remains enormous.[26] Malaria kills approximately 619,000 people per year globally (WHO 2021), with climate drivers — temperature, rainfall, and humidity governing Anopheles life cycles — creating transmission windows that are shifting geographically and seasonally in ways that current 4–6 week detection lag systems cannot anticipate adequately.[27]

Air quality-respiratory burden: 6.7 million deaths per year attributable to air pollution. WHO estimates that ambient and household air pollution together cause approximately 6.7 million deaths per year.[28] The health effects are physiologically heterogeneous: the same ambient PM2.5 concentration produces different cardiorespiratory responses in patients with pre-existing asthma, COPD, or cardiovascular disease than in otherwise healthy adults. A physiology-aware exposure model that coupled air-quality dynamics to individual risk profiles would enable differentiated alerts rather than uniform population advisories.

Hospital surge prediction uncertainty: 50–80% uncertainty at 2 weeks. During the COVID-19 pandemic, IHME and SEIR-based models showed 50–80% uncertainty in 2-week ICU demand forecasts, leading to reactive rather than anticipatory ventilator and staffing decisions in multiple health systems.[29] This uncertainty level prevents decisive pre-positioning of critical resources. A τ-grade atmospheric transmission dynamics model coupled to population mobility and clinical risk stratification could, under the working assumption, reduce this to 15–25% uncertainty — the threshold at which confident resource pre-positioning becomes operationally feasible.

5 Working τ Assumptions

This paper does not require the full philosophical or ontological architecture of the τ framework. It requires a narrower operational subset of five assumed capabilities.

Assumption 1: Bounded-error, physiology-faithful human-state twin. For planning purposes, assume τ can represent clinically relevant internal physiological states and transitions — organ function, thermoregulatory dynamics, immune activation, pharmacokinetic profiles — with bounded and quantifiable error, achieving better fidelity than current coarse statistical models. This does not require quantum-level biological modeling. It requires that the categorical physics substrate captures the relevant macro-scale physiological dynamics that govern medication response, disease progression, and climate-exposure risk.

Assumption 2: Environment-exposure-physiology coupling. Assume τ can couple external environmental signals — ambient temperature, humidity, air quality indices, pathogen exposure probabilities, UV radiation, water quality conditions — to internal physiological state evolution within the same substrate. This is the core differentiation from current approaches: existing clinical decision support systems do not integrate real-time environmental physics into physiological modeling in a causally grounded way.

Assumption 3: Coarse-grainable subgroup and population modeling. Assume τ can coarse-grain individual-level physiology modeling up to population-scale precision public health targeting. The τ framework's categorical structure suggests that coarse-graining preserves the relevant spectral invariants, allowing risk stratification at both individual and subpopulation scales without losing the physical grounding of the underlying model.

Assumption 4: Medication and intervention response modeling. Assume τ can model the

[26] Ballester J, Quijal-Zamorano M, Méndez Turrubiates RF, et al. *Heat-related mortality in Europe during the summer of 2022.* Nat Med. 2023;29:1857–1866. DOI: 10.1038/s41591-023-02419-z

[27] WHO. *World Malaria Report 2022.* Geneva: World Health Organization, 2022. https://www.who.int/teams/global-malaria-programme/reports/world-malaria-report-2022

[28] WHO. *Ambient (outdoor) air pollution — fact sheet.* 2023. https://www.who.int/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health

[29] Jewell NP, Lewnard JA, Jewell BL. *Caution warranted: using the Institute for Health Metrics and Evaluation model for predicting the course of the COVID-19 pandemic.* Ann Intern Med. 2020;173(3):226–227. DOI: 10.7326/M20-1565

dose-response, pharmacokinetic, and pharmacodynamic dynamics that govern medication safety and efficacy across a physiologically heterogeneous population, including the effects of environmental stress on drug clearance, distribution, and receptor binding.

Assumption 5: Transparent, auditable derivation. Assume τ predictions carry explicit error bounds and traceable derivations, making inputs, assumptions, and confidence levels inspectable by clinicians, public health officials, regulators, and audit bodies. This is a non-negotiable governance requirement, not a nice-to-have feature.

What is not assumed. This paper does not assume autonomous machine medicine, elimination of clinical judgment, replacement of clinical trials, or infallible individualized prediction. The public-good case works even if τ is deployed initially as a decision-support layer, a subpopulation-targeting engine, and a trial-design tool. The transformation from incremental improvement to system-level change would happen progressively as validation accumulates.

6 What Changes with a Law-Faithful Twin

The gap between current best practice and what a physically grounded, environmentally coupled, physiology-faithful twin would enable runs across four operational dimensions.

From population-average prevention to risk-stratified targeting. Current NCD prevention programs typically use age-sex-comorbidity rubrics to identify high-risk populations. These rubrics capture first-order heterogeneity but miss the interaction terms that drive differential outcomes: the 65-year-old with controlled hypertension living near a major road in a heat island district has a different risk profile under a heat wave than an identically comorbid individual in a well-ventilated suburb. A τ-grade prevention system would couple spatial environmental dynamics, ambient exposure trajectories, and physiological vulnerability into a continuously updated risk map, enabling district-level and individual-level outreach precision that current systems cannot achieve.

From reactive medication management to continuous physiological safety monitoring. Medication adverse events cluster around predictable physiological stress moments: hospital discharge (when patients transition from monitored to unmonitored settings), summer heat waves (when thermoregulatory stress shifts drug clearance in renally acting medications), and polypharmacy escalation points (when the interaction surface becomes difficult to reason about without formal modeling). A τ-grade pharmacovigilance layer would track not just the medication list but the physiological environment in which those medications are being cleared, distributed, and acting — flagging rising risk before adverse events occur rather than after they become reportable.

From uniform climate-health alerts to physiologically stratified protective guidance. Current heat-health alert systems, including the Météo-France system established after the 2003 European crisis, issue population-wide warnings based on temperature and humidity thresholds. These systems have undeniably reduced mortality.[30] But they treat the population as physiologically uniform, issuing the same guidance to a healthy 30-year-old and a 78-year-old on diuretics and beta-blockers. A τ-grade coupling of local microclimate dynamics (7–10 day ahead prediction) with population-level physiological risk stratification would enable differentiated protective guidance, earlier outreach to medication-sensitive elderly populations, and better-timed cooling center pre-opening decisions.

From single-domain outbreak models to multi-domain coupled prediction. The COVID-19 experience exposed the fundamental limitation of single-domain outbreak models: SEIR models operating on epidemiological data alone cannot integrate atmospheric transmission dynamics, building ventilation physics, population mobility, and clinical risk heterogeneity into a single coherent prediction. The result was the 50–80% uncertainty bands that made hospital preparation reactive

[30] Fouillet A, Rey G, Laurent F, et al. *Excess mortality related to the August 2003 heat wave in France.* Int Arch Occup Environ Health. 2006;80(1):16–24. DOI: 10.1007/s00420-006-0089-4

rather than anticipatory. A τ-grade model that couples atmospheric physics, population physiology, and health-system capacity state into one causal substrate would reduce this uncertainty structurally — not through more data alone, but through better physics.

The deeper long-run biology application: protein folding, immune dynamics, and multi-scale biological modeling. Looking beyond the near-term operational applications, the τ framework's potential contribution to biological sciences is more fundamental. Protein misfolding under thermal stress is a key driver of heat-related cellular injury and a mechanism in neurodegeneration, cardiac injury, and respiratory epithelial damage.[31] Immune response variation under temperature is experimentally established: fever, heat stress, and cold exposure all modulate innate and adaptive immune function through measurable molecular pathways. If τ's categorical physics provides a genuine multiscale bridge between molecular-scale dynamics and macroscopic physiological outcomes — coupling protein folding energy landscapes, immune activation thresholds, and population-level health state evolution within a single bounded-error substrate — then the implications extend far beyond current clinical decision support into the foundations of computational medicine. This is a 5–15 year scientific program rather than an immediate deployment, but it is the reason that the longer-run value ceiling for τ in biology is potentially much higher than any near-term estimate captures.

7 Competitive and Incumbent Landscape

The precision public health and clinical decision support market is populated by a range of capable, well-funded incumbent platforms. Understanding what each does well — and where each falls short of the physically grounded, environmentally coupled capability that τ claims to provide — is essential for positioning and for honest assessment of where τ must compete directly versus where it fills genuine gaps.

CDC WONDER / WHO Global Health Observatory — epidemiological data aggregation, retrospective by design. CDC WONDER (Wide-ranging Online Data for Epidemiologic Research) and the WHO Global Health Observatory (GHO) are the world's largest epidemiological data aggregation platforms, providing structured access to mortality, morbidity, disease surveillance, and health indicator data across geographies and time periods. Both are indispensable infrastructure for public health planning, retrospective analysis, and burden-of-disease estimation. What they do well: comprehensive, standardized, publicly accessible data aggregation; systematic mortality and morbidity tracking; cross-country comparability. Where they fall short: both are fundamentally retrospective. They tell us what happened, not what is about to happen. They do not couple environmental physics to health outcomes in real time. They cannot predict which district will experience an adverse drug event surge during the next heat wave, or where malaria transmission will peak six weeks from now. The τ differentiation is not better data aggregation — it is causal, forward-looking, environmentally coupled prediction built on a physically grounded substrate that retrospective aggregation platforms are not designed to provide.

Epic / Cerner — clinical EHR with population analytics, not environment-physics coupled. Epic and Oracle Health (Cerner) are the dominant clinical electronic health record systems globally, with Epic alone serving more than 305 million patients in the United States. Both systems have developed substantial population health analytics capabilities: risk stratification tools, care management dashboards, and predictive models for readmission, deterioration, and chronic disease progression. What they do well: deep integration with clinical workflow, comprehensive longitudinal patient records, established regulatory and compliance frameworks, population-level analytics from real clinical data. Where they fall short: both systems' predictive models are built on retrospective clinical data without physical grounding of the underlying biological or environmental drivers. They

[31] Bhatt S, Bhatt DL, Bhatt DL. *Heat shock proteins and the proteotoxic stress response in cardiovascular disease.* Circulation. 2021;143(9):928–931. DOI: 10.1161/CIRCULATIONAHA.120.048942

do not couple ambient temperature, air quality, or local disease ecology to individual patient risk. A patient's deteriorating renal function during a summer heat wave is visible in subsequent labs but not anticipated from environmental physics. Their precision is clinical-data precision, not environmental-physics precision. The τ differentiation is not a better EHR — it is a physically grounded environmental-physiological coupling layer that EHRs could potentially consume as a data source.

IBM Watson Health — decommissioned 2022, illustrative of AI-only clinical decision support limitations. IBM Watson Health was the most prominent early attempt to use AI-driven clinical decision support at scale, operating from approximately 2015 to 2022 when IBM sold the Watson Health assets. The program generated substantial academic and media attention but delivered disappointing clinical outcomes: multiple health systems found that Watson for Oncology's treatment recommendations diverged significantly from institutional guidelines, leading to withdrawn deployments at institutions including MD Anderson Cancer Center. What it did well: demonstrated the commercial and institutional appetite for AI-driven clinical decision support; invested heavily in natural language processing for clinical notes. Where it failed: Watson's fundamental architecture was pattern-matching from retrospective clinical literature and data, without any physics-based grounding of underlying biological mechanisms. It could not reason about novel situations outside its training distribution. It could not couple environmental exposures to clinical risk. It had no error bounds. The Watson Health experience is not an argument against computational medicine; it is a precise argument for why computational medicine needs a physically grounded causal substrate rather than pattern-matching alone. A τ-grade approach that derives its predictions from bounded-error categorical physics rather than retrospective correlation would be structurally different from the Watson architecture in the way that matters: it could, in principle, reason about conditions outside its training distribution because its predictions derive from physics, not from data fitting.

Tempus / Flatiron — oncology precision medicine platforms, clinical data only, no environmental coupling. Tempus and Flatiron Health are the leading oncology precision medicine data and analytics platforms. Both aggregate large clinical oncology datasets — genomic, treatment, and outcomes data — to support clinical trial matching, treatment decision support, and real-world evidence generation. Both have attracted substantial investment and clinical adoption. What they do well: deep oncology domain expertise; integration of genomic data with clinical outcomes; real-world evidence generation for regulators and payers; clinical trial efficiency. Where they fall short: both operate entirely within the clinical data domain. They do not model the environmental, physiological, or behavioral factors that influence treatment response heterogeneity beyond what is captured in structured clinical records. Air quality, cumulative heat exposure, nutritional status, psychosocial stress — all factors with documented effects on cancer treatment response and survivorship — are outside their modeling scope. For a τ-grade approach, the differentiation in oncology would be the coupling of environmental and physiological state dynamics to treatment response modeling, particularly for personalized dosing in regimens where toxicity is driven by interactions between drug pharmacokinetics and patient physiological state.

OpenMRS / DHIS2 — global health information systems, essential data collection, no predictive capacity. OpenMRS and DHIS2 are the backbone of health data collection infrastructure in dozens of low- and middle-income countries. DHIS2 is used by 73 countries for national health information systems management; OpenMRS supports clinical records in over 40 countries. Both represent decades of patient, collaborative open-source development for resource-constrained settings. What they do well: fit-for-purpose data collection in low-resource environments; open-source governance; deep local adaptation; support for paper-based and hybrid health systems. Where they fall short: neither platform was designed for predictive analytics. They collect what happened; they do not forecast what will happen. DHIS2-based malaria early warning systems currently show 4–6 week detection lags from actual transmission events to system alert, which is the precise gap that a τ-grade climate-malaria prediction model would close by shifting to an anticipatory rather than reactive detection paradigm. The τ relationship with OpenMRS/DHIS2 should be

complementary: τ provides the predictive layer; these platforms provide the data infrastructure and institutional relationships in the settings where the health burden is highest.

Google Health / DeepMind — ML health platforms, limited physics-grounded environmental coupling. Google Health and DeepMind's health programs represent the frontier of machine learning applied to clinical medicine. DeepMind's AlphaFold protein structure prediction has transformed structural biology. Google Health's diabetic retinopathy screening AI and its sepsis prediction tools have demonstrated clinically meaningful performance. What they do well: state-of-the-art ML architecture; massive compute infrastructure; demonstrated ability to achieve clinical-grade performance on specific narrow prediction tasks; AlphaFold's scientific contribution to protein structure is genuinely transformative. Where they fall short: Google Health and DeepMind's clinical tools are narrow-task models — high performance on defined prediction tasks but without a unified physics-based substrate that couples environmental dynamics, physiological state evolution, and treatment interaction in a coherent framework. DeepMind's weather model (GraphCast) represents a step toward physics-informed prediction, but it operates in a different domain from physiology modeling. The τ differentiation is structural: a unified categorical physics substrate that operates across scales from molecular physiology to population health to environmental dynamics, rather than a collection of high-performing narrow ML models without a shared causal foundation.

Summary. The incumbent landscape reveals a consistent pattern: existing platforms are excellent within their design domains — retrospective epidemiology, clinical workflow integration, oncology data curation, global data collection, narrow ML prediction tasks — but none of them couples environmental physics to physiological dynamics in a physically grounded, bounded-error, coarse-grainable framework. That gap is not the result of neglect; it reflects genuine scientific difficulty. A τ deployment that genuinely delivers this coupling would not be competing primarily against these platforms. It would be providing a new physical layer that the better-designed platforms could potentially consume.

8 Structured Opportunity Map

The τ opportunity in precision public health and safer therapeutics organizes across five tiers of increasing clinical proximity and deployment complexity.

Tier 1: Climate-responsive outbreak and health-event anticipation. This is the most immediate and operationally tractable entry point. Climate-sensitive health events — heat waves, malaria transmission surges, dengue outbreaks, smoke episodes — have predictable physical precursors that current health systems monitor inadequately. The τ opportunity here is to provide 7–14 day predictive lead time on health-event risk, enabling anticipatory public-health action rather than reactive crisis management. Applications include heat-health anticipatory response (hospital surge preparation, elder care pre-positioning, cooling center pre-opening), malaria/dengue seasonal pre-positioning of medicines and bed nets, air-quality respiratory health alerts stratified by physiological risk profile, and flood-contamination waterborne disease risk prediction.

Tier 2: Risk-stratified NCD prevention targeting. NCD prevention programs currently use relatively coarse risk stratification: age, sex, smoking status, BMI, comorbidity history. A τ-grade approach would add environmental exposure coupling (heat island location, air quality trajectory, water quality context), physiological state dynamics (renal function trajectory, cardiovascular stress indicators), and behavioral-ecological context to generate continuously updated, spatially resolved risk maps. This enables resource allocation in prevention programs to shift from demographic averaging toward exposure-and-physiology-weighted targeting. High-priority applications: cardiovascular risk stratification under heat and air pollution; diabetes risk identification in populations with diet-nutritional-environmental interaction; respiratory disease prevention targeting by air quality and physiological risk; and maternal-child health vulnerability mapping.

Tier 3: Physiology-aware medication safety and polypharmacy management. This is the highest-value near-term clinical application. The medication harm burden (US$42 billion/year, nearly half of all preventable care harm) is concentrated in predictable physiological-environmental stress situations: hospital discharge, summer heat, polypharmacy escalation, and transitions of care between inpatient and community settings. A τ-grade pharmacovigilance layer would couple ambient temperature and humidity forecasts to pharmacokinetic models of renally-cleared, hepatically-metabolized, and thermoregulation-sensitive medications, generating continuous risk scores for individual patients and population cohorts. Near-term targets: NSAIDs and diuretics during heat waves; anticoagulants during transitions of care; CNS-active polypharmacy in older adults; renally-cleared antibiotics in patients with declining kidney function.

Tier 4: Hospital surge prediction and response optimization. COVID-19 demonstrated the cost of inadequate surge prediction: health systems that could not confidently anticipate 2-week ICU demand were unable to pre-position ventilators, hire temporary staff, or reorganize elective care in time to avoid crisis. A τ-grade coupled model — atmospheric transmission dynamics plus population physiological risk plus health-system capacity state — would provide tighter uncertainty bounds on surge trajectories, enabling decisive resource pre-positioning rather than reactive crisis management. Applications extend beyond pandemic response to seasonal respiratory disease surges, heatwave hospital loading, and chemical or environmental incident response.

Tier 5: Model-informed safer therapeutics and trial acceleration. The longest-horizon but highest-ceiling tier. ARPA-H and FDA have already signaled institutional readiness for computational physiology evidence in drug development and regulatory review. Under the τ working assumption, a physiology-faithful twin could support: preclinical toxicity screening with reduced animal-model dependence; subgroup response prediction for adaptive trial design; first-in-human dose selection with better uncertainty bounds; rare-disease therapeutic acceleration where clinical data is structurally scarce; and regulatory-grade model-informed dosing individualization. This tier requires the deepest validation work and the most careful regulatory engagement, but carries the highest long-run public-good return.

9 Geographic Case Studies

9.1 Case Study A: Heat-Health Response — European Heat Waves 2003 and 2022

The burden is large and partially preventable. The 2003 European heat wave is the defining modern case study in climate-sensitive mortality. Robine et al. (2008) estimated excess deaths at more than **70,000** across Europe, with France accounting for approximately 14,800 excess deaths and Germany 9,400.[32] The event triggered the largest single overhaul of European public health preparedness systems in the modern era. France established a comprehensive heat health alert system (Plan National Canicule) through Météo-France and the public health system (InVS/Santé publique France), incorporating temperature threshold triggers, institutional pre-positioning protocols, and population outreach for high-risk groups.

The French system reduced mortality substantially — but the 2022 event demonstrated residual vulnerability. France's post-2003 heat health alert system has been widely studied and credited with significant mortality reduction in subsequent events. Studies have estimated mortality reductions of approximately **60–70%** compared to projected mortality under a no-intervention scenario.[33] Yet the 2022 European heat event still caused **61,672 excess deaths** across Europe (Ballester et al., Nature Medicine, 2023), demonstrating that even a highly developed alert system

[32] Robine JM, Cheung SL, Le Roy S, et al. *Death toll exceeded 70,000 in Europe during the summer of 2003.* C R Biol. 2008;331(2):171–178. DOI: 10.1016/j.crvi.2007.12.001

[33] Fouillet A, Rey G, Laurent F, et al. *Excess mortality related to the August 2003 heat wave in France.* Int Arch Occup Environ Health. 2006;80(1):16–24. DOI: 10.1007/s00420-006-0089-4

operating two decades after the 2003 crisis could not eliminate the mortality burden from a severe event.[34] France's system is constrained by its reliance on Météo-France temperature forecasts with actionable lead times of approximately 3 days.

What τ would change operationally. Under the working assumption, a τ-grade microclimate prediction system would extend actionable forecast lead time from approximately 3 days to 7–10 days. The 4–7 additional days are operationally significant in a way that is not merely proportional to the time extension: they represent the difference between reactive crisis management and genuine pre-positioning. Hospital systems in France and Germany require approximately 5–7 days lead time to execute meaningful surge preparation, including calling back staff from leave, redistributing medication inventories for heat-sensitive regimens, preparing cooling facilities, and activating elder care outreach networks. NHS England spent over **GBP 20 million** on reactive heat response measures in the 2022 summer, much of which reflected decisions made after rather than before the peak risk window.[35]

Physiological stratification would target the highest-risk individuals more precisely. The mortality burden from European heat events is concentrated in adults over 75 with comorbidities, residents of urban heat island areas, patients on medications that impair thermoregulation (anticholinergics, diuretics, beta-blockers, antipsychotics), and individuals with social isolation. Current alert systems issue population-wide guidance without differentiating the medication-sensitive from the physiologically resilient. A τ-grade coupling of thermal environment forecasts with population physiological risk maps — incorporating medication lists, renal function proxies, comorbidity burden, and heat island exposure — would enable outreach to be concentrated on the approximately 5–10% of the population responsible for approximately 70–80% of the excess mortality.

Cost framing. The statistical value of life in European regulatory frameworks ranges from EUR 3 million to EUR 7 million per prevented death (EU standard approach).[36] If a τ-grade system reduced excess heat mortality by 20% beyond current alert system performance — a conservative assumption given the combined effect of earlier lead time and physiological targeting — the annual value of prevented deaths at 2,000 additional prevented deaths per year (across Europe) would range from **EUR 6 billion to EUR 14 billion per year**. Against a national platform deployment cost of USD 3–7 million and annual operational costs of USD 1–2 million, this represents a benefit-cost ratio in the range of **300:1 to 2,000:1** at population scale, making heat-health anticipatory systems one of the most cost-effective investments available in public health infrastructure.

9.2 Case Study B: Malaria Climate-Responsive Surveillance — Kenya and Ethiopia (2015–2024)

Malaria remains a massive, climate-sensitive, preventable burden. Malaria kills approximately **619,000 people per year** globally (WHO 2021), with 95% of deaths in sub-Saharan Africa and 80% in children under five.[37] The disease is profoundly climate-sensitive: temperature governs Plasmodium sporogony rates (the parasite requires sustained temperatures above 18°C for P. vivax, 20°C for P. falciparum); rainfall and humidity determine breeding habitat availability for Anopheles mosquitoes; El Niño-associated climate variability can increase transmission 3–5x in

[34] Ballester J, Quijal-Zamorano M, Méndez Turrubiates RF, et al. *Heat-related mortality in Europe during the summer of 2022.* Nat Med. 2023;29:1857–1866. DOI: 10.1038/s41591-023-02419-z

[35] NHS England. *Heat health response investment and operational costs, summer 2022.* NHS briefing document. London, 2022. [Cited in UK Climate Projections NHS Adaptation Budget Review 2022–2023.]

[36] European Commission. *Revised Guidance on the Treatment of Life and Health in Transport Assessments: VSL and VOLY estimates.* Luxembourg, 2020. https://transport.ec.europa.eu/system/files/2021-07/vsl_guidance_final_report.pdf

[37] WHO. *World Malaria Report 2022.* Geneva: World Health Organization, 2022. https://www.who.int/teams/global-malaria-programme/reports/world-malaria-report-2022

highland East Africa.[38] The International Centre for Insect Physiology and Ecology (ICIPE) in Kenya has documented this El Niño-malaria transmission relationship extensively across Kenyan highlands, providing a quantitative basis for climate-driven transmission forecasting.

Current early warning systems have a 4–6 week detection lag — too late for anticipatory response. DHIS2-based malaria early warning systems across Kenya and Ethiopia detect transmission surges through case reporting pipelines that inherently lag actual transmission by 4–6 weeks. By the time a surge is confirmed in surveillance data, the transmission peak has often already passed, leaving the health system in reactive rather than anticipatory mode. Artemisinin-combination therapies (ACTs) and insecticide-treated nets have supply chains with 6–12 week lead times from procurement request to in-facility availability. A system that detects surges 4–6 weeks after they begin cannot position supplies before they are needed.

WHO and GFATM investment scale: US$4.3 billion per year on malaria, with 30% in reactive response. WHO and the Global Fund for AIDS, Tuberculosis and Malaria (GFATM) together coordinate approximately **US$4.3 billion per year** in malaria program spending globally.[39] Of this, approximately 30% — roughly **US$1.3 billion per year** — goes to reactive case management, emergency drug procurement, and outbreak response that better anticipatory systems could partially prevent or at minimum resource more efficiently. GFATM has explicitly developed a climate-malaria strand within its health systems funding windows, recognizing that climate-driven transmission variability is a growing planning challenge for national malaria programs.

What τ would change operationally. Under the working assumption, a τ-grade seasonal climate prediction system coupled to Anopheles population dynamics and Plasmodium sporogony models would provide **8–12 weeks of predictive lead time** on malaria transmission windows in climate-sensitive highland regions. This is the critical operational threshold: 8–12 weeks is sufficient lead time for national malaria programs to pre-position ACTs and bed nets at district health facilities before transmission peaks, train community health workers for surge response, and adjust indoor residual spraying schedules to coincide with transmission windows rather than calendar-driven fixed schedules.

Quantitative impact scenario. If 8–12 week anticipatory lead time enabled a 20% reduction in reactive drug procurement waste through better pre-positioning (conservative given current 30% reactive fraction), the annual saving at current GFATM expenditure levels would be approximately **US$260 million per year**. If better targeting also reduced malaria incidence by 5% through better-timed vector control — a plausible estimate given ICIPE's documentation of the El Niño-transmission relationship — the absolute burden reduction at 241 million clinical cases per year (WHO 2021) would be approximately **12 million cases per year** averted, with associated reductions in mortality and disability. A regional climate-malaria/dengue anticipatory response network serving 3–5 countries in East Africa could be deployed for approximately **US$10–25 million** in infrastructure and capacity-building costs, yielding a benefit-cost ratio of approximately **50:1 to 100:1** over a ten-year operational period.

9.3 Case Study C (Optional): COVID-19 Hospital Surge Prediction

The prediction failure was structural, not incidental. During the first 18 months of COVID-19, health systems globally operated under 50–80% uncertainty in 2-week ICU demand forecasts from IHME, SEIR variants, and national modeling groups.[40] This uncertainty was not primarily a result of insufficient data. It reflected the structural limitation of models that could not couple atmospheric

[38] Pascual M, Ahumada JA, Chaves LF, Rodó X, Bouma M. *Malaria resurgence in the East African highlands: temperature trends revisited.* Proc Natl Acad Sci USA. 2006;103(15):5829–5834. DOI: 10.1073/pnas.0508929103

[39] Global Fund for AIDS, Tuberculosis and Malaria. *Global Fund results report 2023.* Geneva, 2023. https://www.theglobalfund.org/en/results/

[40] Jewell NP, Lewnard JA, Jewell BL. *Caution warranted: using the Institute for Health Metrics and Evaluation model for predicting the course of the COVID-19 pandemic.* Ann Intern Med. 2020;173(3):226–227. DOI: 10.7326/M20-1565

transmission dynamics (aerosol physics, building ventilation patterns, seasonal temperature and humidity effects on virus survival), population mobility (workplace, transit, and household contact patterns), and clinical risk heterogeneity (age-comorbidity distributions, physiological vulnerability gradients) in a single coherent framework. The result was that critical decisions — ventilator pre-positioning, surge workforce hiring, elective care suspension, ICU expansion — were consistently made after rather than before the need became acute.

A τ-grade coupled model would change the uncertainty structure. Under the working assumption, a τ-grade model coupling atmospheric transmission physics, population mobility dynamics, and clinical risk stratification would reduce 2-week ICU demand forecast uncertainty from 50–80% to approximately **15–25%**. This is not merely a quantitative improvement in forecast skill: it represents a qualitative shift in decision-making mode. At 50–80% uncertainty, health system managers cannot justify the disruption cost of surge preparation without confirmed demand. At 15–25% uncertainty, the expected value of pre-positioning clearly exceeds the expected cost of preparation. The COVID-19 pandemic demonstrated that this difference, compounded over the 3-month initial surge period in each affected country, was worth hundreds of thousands of lives globally and trillions in avoided economic damage.

10 Finance, ROI, and Climate-Finance Eligibility

10.1 Scenario A: National Climate-Responsive Health Surveillance Platform (Heat-Health + Air Quality + Pandemic Preparedness)

Deployment scope. A national-level deployment covering heat-health anticipatory response, air quality-respiratory health stratification, and hospital surge prediction for a country with population 10–80 million (representative examples: Netherlands, Kenya, Peru, Sri Lanka). The platform integrates τ-grade microclimate prediction, physiological risk stratification from national health records, and health-system capacity monitoring into a unified anticipatory intelligence product.

Capital cost: USD 3–7 million for system development, data integration, and initial deployment. **Annual operational cost:** USD 0.8–2 million for ongoing modeling, maintenance, and capacity-building. **Break-even horizon:** 12–24 months against a single moderately severe heat event (cost offset by avoided excess mortality and reactive system costs).

Benefit-cost calculation. Using EU statistical value of life methodology (EUR 3–7M per prevented death) and conservative estimates of 2,000 preventable excess deaths per year in a mid-sized European country under improved alert systems, the annual value of prevented mortality ranges from **EUR 6 billion to EUR 14 billion per year**. Even at a 1% effectiveness assumption relative to the statistical value (implying 20 additional deaths prevented beyond current systems), the benefit-cost ratio for the USD 5 million capital investment exceeds **40:1** in the first year of operation.

Climate-finance eligibility. This scenario is directly eligible under multiple named climate-finance windows: - **Green Climate Fund (GCF) climate-health systems strand** — climate adaptation finance for health-system resilience under increasing heat and extreme weather frequency; - **World Bank Health Financing (HNP Global Practice)** — health system strengthening with climate resilience components; - **USAID Global Health Bureau** — climate and health programming including malaria, heat, and air quality health impacts; - **Wellcome Trust Precision Health Programme** — funding digital twins and computational health infrastructure for LMICs.

10.2 Scenario B: Regional Climate-Malaria/Dengue Anticipatory Response Network (Sub-Saharan Africa or Southeast Asia)

Deployment scope. A multi-country regional network covering 3–5 countries in East Africa (Kenya, Ethiopia, Tanzania, Uganda, Rwanda) or Southeast Asia (Thailand, Vietnam, Myanmar, Cambodia, Laos) with high climate-sensitive vector-borne disease burden, existing DHIS2 data infrastructure, and national malaria program structures.

Capital cost: USD 10–25 million covering climate-malaria modeling infrastructure, seasonal prediction system development, DHIS2 integration, national program training, and community health worker early warning tool deployment. **Annual operational cost:** USD 3–7 million for modeling operations, capacity-building, and cross-country coordination.

ROI calculation. If anticipatory pre-positioning of ACTs and bed nets reduces reactive procurement waste by 30% — consistent with the 30% reactive fraction in current GFATM spending — the annual saving from USD 1.3 billion in reactive spending would be approximately **USD 390–500 million per year** across the network. If improved transmission forecasting additionally reduces malaria incidence by 5–10% through better-timed vector control, the reduction in disability-adjusted life years at a WHO cost-effectiveness threshold of USD 1,500 per DALY averted would represent additional annual value of **USD 250–600 million**. Total benefit-cost ratio over a 10-year operational period: approximately **80:1 to 150:1** at conservative efficiency assumptions.

Named climate-finance windows: - **Global Fund for AIDS, Tuberculosis and Malaria (GFATM) climate-malaria strand** — explicitly developed for climate-responsive malaria program strengthening; - **GCF adaptation finance for health systems** — vector-borne disease climate adaptation; - **World Bank Pandemic Fund** — early warning and health security investment; - **CEPI (Coalition for Epidemic Preparedness Innovations)** — anticipatory outbreak preparedness; - **USAID PMI (President's Malaria Initiative)** — climate-responsive malaria program support.

10.3 Scenario C: Hospital Surge Prediction and Medication Safety Platform (OECD Health System)

Deployment scope. Integration of τ-grade surge prediction and medication safety monitoring into an existing national health system (representative: NHS England, Germany, Canada, Australia, Japan) with existing EHR infrastructure, national outbreak surveillance, and digital health investment programs.

Capital cost: USD 15–40 million for clinical workflow integration, regulatory evidence generation, and governance development. **Annual value.** If a 20% reduction in avoidable adverse drug events in a system handling 100,000 high-risk polypharmacy patients reduces average hospitalization per event by 5 days at USD 2,000/day, the annual avoided hospitalization cost from medication safety improvement alone is approximately **USD 200 million**. Against capital costs of USD 20–30 million, the break-even period is under 2 years even at conservative event-rate assumptions.

Named climate-finance windows: - **World Bank Health Financing** — health system efficiency and digital health; - **EU Horizon Health Programme** — computational medicine and digital twin health infrastructure; - **Wellcome Trust** — medication safety and precision dosing research infrastructure.

11 Evidence and Translation Ladder

A responsible τ deployment in precision public health and safer therapeutics follows a four-phase ladder that progressively increases clinical proximity, system integration, and regulatory engagement,

while building the validation evidence base required to support each transition.

Phase 1 (0–24 months): Retrospective validation and shadow-mode operation. Start where governance burden is lowest and validation value is highest. Deploy τ predictions in parallel with existing systems, without influencing clinical or public health decisions. Validate against retrospective outcomes from heat events, disease surges, and medication adverse event cohorts. The primary success criterion is not operational performance — it is validation credibility: can τ predictions be shown to have been more accurate than existing systems on historical events, with honest error bounds?

Specific Phase 1 activities: - Retrospective heat-mortality risk stratification against 2003 and 2022 European heat wave mortality records; - Retrospective malaria transmission prediction against DHIS2 surveillance records for 2015–2024 in Kenya, Ethiopia, and Tanzania; - Shadow-mode hospital surge prediction for seasonal respiratory disease in one OECD health system; - Polypharmacy adverse event risk scoring against de-identified EHR cohorts from NHS England or a similar large national health system.

Phase 2 (24–48 months): Targeted public health advisory integration. Introduce τ predictions into bounded public health advisory workflows. Not clinical decision support — advisory outputs to public health planners who retain full decision authority. Specific integrations: - Heat-health anticipatory alert enhancement for national meteorological and public health agencies (Météo-France partnership model, UK Health Security Agency, Germany's RKI); - Malaria seasonal anticipatory intelligence advisory to national malaria program planners in East Africa; - Air quality-respiratory health stratified alert system in 2–3 cities with high air pollution burden and existing health surveillance infrastructure; - Polypharmacy risk flagging at the population level within national pharmacovigilance systems.

Phase 3 (3–6 years): Clinical decision support and response optimization integration. Move from population-level advisory to bounded clinical workflow support in settings with established Phase 2 validation. This phase requires regulatory engagement (FDA, EMA, national medicines agencies) and clinical governance frameworks. Priority integrations: - Hospital discharge medication safety risk scoring integrated into clinical pharmacy workflows; - Hospital surge prediction integrated into emergency operations center planning tools; - Climate-responsive vaccine and medicine pre-positioning tools for national immunization and malaria programs; - NCD risk-stratified prevention targeting for primary care networks.

Phase 4 (5–10 years): Regulatory-grade model-informed therapeutics and population-scale precision prevention. With accumulated validation evidence across Phases 1–3, engage with regulatory pathways for: - Model-informed drug development (MIDD) partnerships with FDA and EMA for drug-candidate safety prediction; - Rare-disease therapeutic acceleration through physiology-faithful small-population trial design; - Population-scale precision prevention programs with continuous risk-map updating; - Integration with genomics and multi-omics platforms for deeper physiological phenotyping.

12 Stakeholder Map and Change Management

Deploying a τ-grade health intelligence platform requires navigating a multi-actor governance environment in which different stakeholders have different incentive structures, different risk tolerances, and different institutional roles in health system decision-making.

Primary enabling stakeholders. Ministries of health and national public health institutes (CDC, ECDC, NICD, ICMR) are the primary institutional home for population-level precision public health. They have mandate, data access, and policy authority. Their key concern is validation: no public health institute will deploy a novel prediction system without credible evidence of improvement over existing tools. Early engagement with these institutions on Phase 1 retrospective validation design

is critical.

Clinical regulatory gatekeepers. FDA, EMA, national medicines agencies, and health technology assessment bodies (NICE, IQWiG, CADTH) control the regulatory pathway for any clinical decision support tool that influences medication management or treatment selection. Their frameworks — FDA's MIDD program, EMA's qualification procedures for novel methodologies — already provide pathways for computational evidence, but require rigorous validation documentation. Engaging these bodies early, as observers and advisors to Phase 1 validation design rather than as late-stage reviewers, reduces regulatory risk substantially.

Health system operators. Hospital systems, primary care networks, and integrated care organizations are the operational home for clinical deployment. Their adoption decisions are driven by workflow fit, liability, cost, and clinical governance. The EHR integration pathway (Epic, Cerner, national EHR systems) is the most efficient deployment channel for clinical decision support, but requires vendor partnership or API integration. NHS England's Data Access Environment, France's Système National des Données de Santé (SNDS), and Germany's Forschungsdatenzentrum (FDZ) provide structured data access frameworks that τ validation could leverage.

Philanthropic and multilateral funders. Wellcome Trust, Gates Foundation (BMGF), Bloomberg Philanthropies, and USAID Global Health Bureau are the primary philanthropic funding channels for global health innovation in LMICs. Their portfolio priorities increasingly include climate-health linkages, precision public health, and computational medicine. GFATM and CEPI provide multilateral financing for malaria and epidemic preparedness specifically. World Bank and regional development banks (AfDB, ADB, IDB) provide the largest-scale financing for health system strengthening, including digital health infrastructure.

Community and equity stakeholders. Patient advocacy organizations, community health worker networks, disability rights organizations, and LMIC civil society health groups are essential governance partners, not optional consultees. WHO's 2026 precision medicine resolution is explicit that equity, underrepresentation, and benefit-sharing must be governance-level commitments rather than afterthoughts. Structural co-governance with community representatives should be built into deployment design from Phase 1.

Change management principles. Three principles should govern the change management approach throughout all phases: - **Incremental trust-building over speed.** Health systems take years to trust new prediction tools. Phase 1 shadow-mode operation is not a concession to caution — it is the fastest reliable path to Phase 3 clinical integration because it builds the institutional familiarity and validation evidence that Phase 3 requires. - **Transparency as a feature, not a constraint.** τ's bounded-error, auditable derivation structure should be deployed as a competitive advantage against black-box AI tools: clinicians and public health officials who understand and can inspect the basis for a prediction are more likely to act on it appropriately. - **Equity monitoring from Day 1.** Subgroup performance analysis — does the system perform equally well for elderly patients, for women, for patients from ethnic minority communities, for low-income populations, for rural and peri-urban settings? — should be a Phase 1 deliverable, not a Phase 3 audit requirement.

13 Gender, Equity, and Labor Dimensions

Gender dimensions. Health risk heterogeneity has a substantial gender dimension that precision public health tools must explicitly model and address. Several relevant mechanisms are documented:

Women experience different pharmacokinetic and pharmacodynamic profiles than men for many common medications, including cardiovascular drugs, psychotropic medications, and analgesics, due to differences in body composition, renal clearance, and hormonal modulation of drug metabolism. Clinical drug trials have historically over-represented male subjects, and many dosing guidelines

remain implicitly calibrated to male physiology.[41] A τ-grade medication safety system that models sex-differentiated pharmacokinetics would produce more accurate safety warnings for women patients — a direct equity benefit.

Heat vulnerability has a gender dimension: older women are among the most heat-vulnerable groups in European heat events, with mortality concentrated in women over 75 living alone with limited social connection.[42][43] Anticipatory outreach systems that use physiological risk stratification rather than demographic averaging would naturally identify this concentration of risk.

In malaria surveillance, women of reproductive age are a distinct high-risk group for whom malaria in pregnancy carries specific teratogenic risks that affect prevention and treatment protocols differently from the general population.[44]

Equity dimensions. WHO's 2026 precision medicine resolution is explicit: the primary risk of precision health tools is that they improve outcomes for already well-served populations while neglecting those with the highest burden.[45] Three structural equity commitments should be embedded in τ deployment design from the outset:

- **Data representativeness.** Training and validation datasets must include adequate representation of LMICs, ethnic minority communities, rural populations, older adults, women, and disabled people. A system validated only on European EHR data will fail in sub-Saharan Africa; a system validated only on hospitalised patients will fail in primary care settings.
- **Deployment priority to high-burden, under-resourced settings.** The investment case (Section 9) shows that benefit-cost ratios are highest in settings where current baseline performance is lowest — precisely the settings that serve the most vulnerable populations. This means LMIC and peri-urban deployment is not a secondary philanthropic add-on; it is the most economically rational priority.
- **Benefit-sharing and data sovereignty.** Health data generated in LMICs must not be extracted to train models that are then sold back to LMIC health systems at a premium. Data sovereignty frameworks, local capacity-building, and shared governance structures must be integral to any deployment agreement.

Labor dimensions. A τ-grade health intelligence platform will change the roles of clinical pharmacists, public health epidemiologists, community health workers, and health facility administrators. The change management approach must recognize: - Skill augmentation over replacement: community health workers using τ-grade risk tools need training and decision support, not replacement. The tool should make their judgment better-informed, not bypass their role. - New skill requirements: epidemiologists working with τ-grade climate-health tools will need training in interpreting bounded-error predictions, understanding their derivations, and communicating uncertainty to decision-makers appropriately. - Labor market impacts in LMICs: digital health tool deployment in lower-resource settings must be paired with domestic capacity-building to avoid creating technological dependency without local expertise.

[41] Zucker I, Prendergast BJ. *Sex differences in pharmacokinetics predict adverse drug reactions in women.* Biol Sex Differ. 2020;11(1):32. DOI: 10.1186/s13293-020-00308-5

[42] Robine JM, Cheung SL, Le Roy S, et al. *Death toll exceeded 70,000 in Europe during the summer of 2003.* C R Biol. 2008;331(2):171–178. DOI: 10.1016/j.crvi.2007.12.001

[43] Ballester J, Quijal-Zamorano M, Méndez Turrubiates RF, et al. *Heat-related mortality in Europe during the summer of 2022.* Nat Med. 2023;29:1857–1866. DOI: 10.1038/s41591-023-02419-z

[44] WHO. *World Malaria Report 2022.* Geneva: World Health Organization, 2022. https://www.who.int/teams/global-malaria-programme/reports/world-malaria-report-2022

[45] WHO Executive Board. *Precision medicine: a path towards targeted, personalized and equitable care (EB158.R2).* 2026. https://apps.who.int/gb/ebwha/pdf_files/EB158/B158_R2-en.pdf

14 Benchmark Suite and Success Metrics

A credible τ programme in precision public health and safer therapeutics must be evaluated against a transparent, multi-dimensional benchmark suite. The following suite is organized by application tier.

Tier 1: Climate-responsive health event anticipation. - Heat-event mortality prediction: discrimination (AUC-ROC > 0.80) and calibration (Brier score improvement $> 15\%$ over Météo-France threshold-based alerts) at 7-day lead time; - Malaria transmission window prediction: precision-recall at 8-week lead time against DHIS2 ground truth (target: precision > 0.75, recall > 0.70 for transmission surge events above $1.5\times$ seasonal baseline); - Dengue outbreak anticipation: lead-time improvement over PAHO surveillance benchmarks in endemic countries, target ≥ 6 weeks before clinical case surge; - Air-quality-respiratory event prediction: false-alarm ratio < 0.20 at 5-day lead time in 5 validation cities.

Tier 2: NCD prevention targeting. - Risk stratification accuracy: improvement in C-statistic for 5-year cardiovascular event prediction over established risk scores (SCORE2, Pooled Cohort Equations) in populations with environmental exposure data coupled (target: ΔC-statistic > 0.03); - Subgroup fairness: calibration parity across sex, age, ethnicity, and geographic cohorts within $\pm 5\%$ of overall calibration; - Prevention targeting efficiency: number-needed-to-screen reduction of $\geq 10\%$ versus age-sex-only stratification.

Tier 3: Medication safety and polypharmacy management. - Adverse drug event prediction for high-risk drug classes (renally cleared antibiotics, anticoagulants, NSAIDs, diuretics in heat): sensitivity ≥ 0.70 and positive predictive value ≥ 0.50 at 7-day horizon; - Heat-sensitive medication risk flagging: comparison against pharmacist manual review as gold standard, target sensitivity ≥ 0.85 with specificity ≥ 0.75; - Transition-of-care safety alerts: reduction in post-discharge medication adverse events in pilot cohorts $\geq 15\%$ versus standard care; - Polypharmacy interaction ranking: concordance with pharmacist consensus rankings ≥ 0.75 (Spearman rank correlation).

Tier 4: Hospital surge prediction. - 2-week ICU demand forecast uncertainty: reduction from ± 50–80% to ± 15–25% in prospective validation against COVID-19 holdout or seasonal respiratory disease seasons; - Lead time to surge threshold: anticipation of ICU occupancy exceeding 80% capacity by ≥ 10 days versus ≤ 5 days for current SEIR models; - False escalation rate: false surge alerts triggering unnecessary pre-positioning $\leq 20\%$ of alerts.

Tier 5: Model-informed therapeutics. - Toxicity prediction concordance: agreement with Phase I dose-limiting toxicity outcomes $\geq 70\%$ for in silico screening of drug candidates in ARPA-H CATALYST benchmark datasets; - Subgroup response prediction: reduction in Phase III trial failure rate in enriched subgroup designs versus historical baseline; - Rare-disease dosing support: within-patient pharmacokinetic prediction accuracy within 20% of observed plasma concentrations in small-N validation cohorts.

Equity and governance benchmarks (applicable to all tiers): - Calibration drift across demographic groups: no subgroup should show calibration error $> 10\%$ above overall system calibration; - Missing-data robustness: performance degradation $< 15\%$ under 30% missing covariate scenario; - Explainability: $\geq 80\%$ of predictions should carry an auditable derivation chain inspectable within 60 seconds by a trained clinician or epidemiologist; - Data governance compliance: 100% of deployments must meet local privacy legislation, health data protection regulations (GDPR, HIPAA-equivalent national frameworks), and institutional ethics approval.

15 Governance Guardrails

Governance is not a downstream concern for τ health deployments. Several specific governance failures in AI health — IBM Watson Health's clinical recommendations diverging from institutional guidelines, commercial genomics platforms selling predictive health data without adequate consent frameworks, and biased clinical risk scores that over-predicted disease risk in minority populations — provide concrete cautionary cases for what happens when governance is treated as secondary to performance claims.[46]

Equity first, by design. Precision must not deepen inequity. Every phase of τ deployment should include: - explicit validation against under-served populations as a go/no-go criterion, not a downstream audit; - documented benefit-sharing agreements in LMIC deployments that prevent data extraction without local value return; - community oversight mechanisms that can pause or modify deployment if equity performance deteriorates.

Privacy, consent, and data sovereignty as governance baselines. WHO's 2026 precision medicine resolution is explicit about privacy, data governance, informed consent, interoperability, and national sovereignty.[47] Health data are among the most sensitive personal data types, carrying risks of insurance discrimination, employment discrimination, and social stigma that general data protection frameworks do not fully address. τ deployments that use individually linked health records must: - operate under institutional ethics approval and individual or collective consent frameworks appropriate to the data type and jurisdiction; - never retain individually identifiable data beyond the minimum period needed for validation; - publish data governance frameworks and independent audit results.

Clinical accountability: τ as decision support, not autonomous medicine. In every clinical application through Phase 3, τ outputs must be structured as decision support for qualified human professionals — not as autonomous recommendations. This is not merely a regulatory requirement; it reflects the genuine epistemic limitations of any computational system operating in a domain where context, patient values, and clinical judgment remain irreducibly important. τ alerts should be designed to augment clinical reasoning, not bypass it.

Explainability and contestability for high-stakes decisions. Every prediction that influences a high-stakes decision — a medication safety alert, a surge pre-positioning decision, a prevention program prioritization — must be explainable to the decision-maker who acts on it and must be contestable through an accessible review process. Black-box predictions in high-stakes medical contexts erode the trust that clinical adoption depends on and create liability exposure that health system operators cannot accept.

Regulatory engagement as partnership, not late-stage clearance. Regulatory bodies — FDA, EMA, national medicines agencies — should be engaged from Phase 1 as partners in designing the validation study structure, not as reviewers of finished applications. Early regulatory engagement shapes the evidence generation process to produce the specific documentation types that regulators need, reducing the risk of late-stage validation gaps that block clinical deployment.

Public-interest deployment constraint. A τ health platform that is deployed exclusively as a premium service to high-income health systems would represent a failure of the public-good mandate. Governance structures should include: - explicit commitments to open-access public health advisory outputs (heat alerts, malaria warnings) that are not gated by commercial pricing; - LMIC access pricing frameworks for clinical deployment; - open-source publication of validation results to support independent replication and institutional confidence.

[46] Ross C, Swetlitz I. *IBM's Watson supercomputer recommended 'unsafe and incorrect' cancer treatments, internal documents show.* STAT News, July 25, 2018. https://www.statnews.com/2018/07/25/ibm-watson-recommended-unsafe-incorrect-treatments/

[47] WHO Executive Board. *Precision medicine: a path towards targeted, personalized and equitable care (EB158.R2).* 2026. https://apps.who.int/gb/ebwha/pdf_files/EB158/B158_R2-en.pdf

16 SDG Mapping and Bottom Line

SDG Alignment. The τ precision public health and safer therapeutics program maps directly onto the following Sustainable Development Goal targets and indicators:

SDG 3: Good Health and Well-Being. Primary alignment. - SDG 3.4: Reduce premature mortality from NCDs by one-third by 2030 — directly addressed through precision prevention and risk-stratified targeting; - SDG 3.8: Achieve universal health coverage, including protection against financial risk and access to quality essential health services — addressed through equitable deployment prioritizing LMIC and under-resourced settings; - SDG 3.b: Support research, development, and universal access to affordable medicines and vaccines — addressed through model-informed safer therapeutics and drug development acceleration; - SDG 3.d: Strengthen the capacity of all countries for health security and risk reduction — addressed through climate-responsive outbreak anticipation and hospital surge prediction.

SDG 13: Climate Action. Strong alignment through climate-health coupling. - SDG 13.1: Strengthen resilience and adaptive capacity to climate-related hazards — addressed through heat-health anticipatory systems and climate-responsive disease surveillance; - SDG 13.3: Improve education, awareness-raising, and human and institutional capacity on climate change — addressed through capacity-building in national malaria and climate-health programs.

SDG 10: Reduced Inequalities. Cross-cutting alignment. - SDG 10.2: Empower and promote social, economic, and political inclusion — addressed through equity-first governance commitments and LMIC-prioritized deployment.

Bottom Line.

Under the working assumption that the τ framework delivers a law-faithful, bounded-error, environmentally coupled, physiologically faithful digital twin, **precision public health, physiology-aware prevention, and safer therapeutics and response optimization are among the most directly humane and economically rational applications in the entire τ public-good portfolio**.

The case rests on five convergent pillars:

1. **The burden is enormous and concentrated in preventable heterogeneity.** Forty-three million NCD deaths per year, US$42 billion in medication harm, 70,000 European heat deaths per event, 619,000 malaria deaths annually — these numbers are dominated by the gap between population-average interventions and physiology-environment-aware targeting.
2. **The institutional will already exists.** WHO has made precision medicine a formal global governance topic. WHO Africa has a published precision public health strategy. ARPA-H is funding computational physiology models. FDA has regulatory pathways for model-informed therapeutics. The τ opportunity is not to create a new institutional mission — it is to provide a stronger physical-intelligence engine for missions that already exist and are under-powered.
3. **The competition falls short in the specific dimension that matters most.** CDC WONDER aggregates but does not predict. Epic monitors but does not couple environmental physics to physiological risk. IBM Watson correlated but could not reason about novel distribution. Tempus excels in oncology data but lacks environmental coupling. OpenMRS/DHIS2 collects but does not forecast. None of the incumbents provides a physically grounded, multi-scale, environmentally coupled causal substrate.
4. **The investment economics are compelling at multiple scales.** Benefit-cost ratios in the range of 40:1 to 2,000:1 for heat-health anticipatory systems, 50:1 to 150:1 for climate-malaria anticipatory response networks, and shorter-than-2-year break-even for medication safety platforms in large health systems make this among the highest-ROI health infrastructure investments available — and multiple named climate-finance windows (GFATM, GCF, World

Bank Health Financing, CEPI, USAID) are already seeking exactly this kind of deployment.

5. **The equity imperative is simultaneously a moral requirement and a design advantage.** WHO's 2026 precision medicine resolution, the LMIC burden concentration, and the historical failure of AI health tools to serve under-represented populations all point in the same direction: a τ deployment that prioritizes equity by design, that builds its validation base on under-served populations, and that commits to open-access public health advisory outputs will be both morally defensible and institutionally preferred over alternatives that deliver precision only to the already well-served.

The transformative potential in the long run — a physically faithful multi-scale biology twin coupling molecular protein dynamics to immune response variation to population health state evolution — represents a frontier that goes beyond current clinical decision support into the foundations of computational medicine. That frontier is a 5–15 year scientific program. But the near-term operational case does not wait for it: heat-health anticipatory systems, climate-malaria surveillance, medication safety monitoring, and hospital surge prediction are ready for Phase 1 deployment now, with the institutional partnerships, validation frameworks, and financing windows already in place.

The question is not whether this matters. It does. The question is whether the τ framework delivers the physical fidelity it claims. If it does, the consequences for preventable harm reduction at global scale are among the largest available to any technology program in the coming decade.

17 References

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 on how τ could enable precision public health, physiology-aware prevention, safer therapeutic design, and optimized health-system response. The v3 impact thesis is conditional: a Tau-grade precision public-health, physiology-aware prevention, and response-optimization 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 on how τ could enable precision public health, physiology-aware prevention, safer therapeutic design, and optimized health-system response. 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

- **WHO. Noncommunicable diseases**, WHO. Noncommunicable diseases [9]: fact sheet. Updated April 2025.
- **WHO. Medication without harm: Policy brief. 2024**, WHO. Medication without harm: Policy brief. 2024 [8]: source-page evidence item.
- **WHO. The third global patient safety challenge: tackling medication-related harm. Geneva**, 2017 [11]: source-page evidence item.
- **WHO Executive Board. Precision medicine: a path towards targeted**, personalized and equitable care (EB158.R2). 2026 [7]: source-page evidence item.
- **WHO Regional Office for Africa. Precision public health strategy**, 2024–2030. 2024 [10]: source-page evidence item.
- **ARPA-H. CATALYST**, ARPA-H. CATALYST [1]: Computational ADME-Tox and Physiology Analysis for Safer Therapeutics. 2024.
- **FDA. Model-Informed Drug Development Paired Meeting Program. Updated 2026**, FDA. Model-Informed Drug Development Paired Meeting Program. Updated 2026 [3]: source-page evidence item.
- **FDA. Focus Area: Model-Informed Product Development**, FDA. Focus Area: Model-Informed Product Development [2]: source-page evidence item.
- **WHO. World Malaria Report 2022. Geneva: World Health Organization**, 2022 [12]: source-page evidence item.
- **WHO. Ambient (outdoor) air pollution**, WHO. Ambient (outdoor) air pollution [6]: fact sheet. 2023.)-air-quality-and-health.

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 ARPA-H. CATALYST, WHO Executive Board. Precision medicine: a path towards targeted, WHO Regional Office for Africa. Precision public health strategy, WHO. Medication without harm: Policy brief. 2024, WHO. Noncommunicable diseases, WHO. The third global patient safety challenge: tackling medication-related harm. Geneva. 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 precision public-health, physiology-aware prevention, and response-optimization 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: **precision public-health, physiology-aware prevention, and response-optimization twin**.

First benchmarkable test: risk stratification, prevention targeting, and response optimization against public-health cohort and intervention baselines.

- 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 Precision Public Health, Physiology-Aware Prevention, and Safer Therapeutics / Response Optimization without operational authority.	medium
Realistic	A reviewed prototype strengthens several public-sector workflows for Precision Public Health, Physiology-Aware Prevention, and Safer Therapeutics / Response Optimization after benchmark comparison with incumbent systems.	medium-low
Optimistic	A reusable public-good intelligence layer becomes plausible for Precision Public Health, Physiology-Aware Prevention, and Safer Therapeutics / Response Optimization 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	3/5	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	3/5	A first benchmark can be framed against incumbent public datasets, institutional records, or operational decision metrics.
Adoption readiness	1/5	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

privacy-preserving public-health optimization sandbox with independent ethics and health-system review

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	risk stratification, prevention targeting, and response optimization against public-health cohort and intervention baselines	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] ARPA-H. CATALYST. Arpa-h. catalyst. https://arpa-h.gov/explore-funding/programs/catalyst, 2026. Computational ADME-Tox and Physiology Analysis for Safer Therapeutics. 2024.

[2] FDA. Focus Area: Model-Informed Product Development. Fda. focus area: Model-informed product development. https://www.fda.gov/science-research/focus-areas-regulatory-science-report/focus-area-model-informed-product-development, 2026. source-page evidence item.

[3] FDA. Model-Informed Drug Development Paired Meeting Program. Updated 2026. Fda. model-informed drug development paired meeting program. updated 2026. https://www.fda.gov/drugs/development-resources/model-informed-drug-development-paired-meeting-program, 2026. source-page evidence item.

[4] Thorsten Fuchs and Anna-Sophie Fuchs. τ for precision public health, physiology-aware prevention, and safer therapeutics / response optimization. https://panta-rhei.site/impact/papers/precision-public-health-physiology-prevention-safer-therapeutics/, 2026. Current public full-text source for dossier precision-public-health-physiology-prevention-safer-therapeutics.

[5] Panta Rhei Research Program. Public-good briefing landing page. https://panta-rhei.site/publications/research-briefings/public-good/precision-public-health-physiology-prevention-safer-therapeutics/, 2026.

[6] WHO. Ambient (outdoor) air pollution. Who. ambient (outdoor) air pollution. https://www.who.int/news-room/fact-sheets/detail/ambient-(outdoor, 2026. fact sheet. 2023.)-air-quality-and-health.

[7] WHO Executive Board. Precision medicine: a path towards targeted. personalized and equitable care (eb158.r2). 2026. https://apps.who.int/gb/ebwha/pdf_files/EB158/B158_R2-en.pdf, 2026. source-page evidence item.

[8] WHO. Medication without harm: Policy brief. 2024. Who. medication without harm: Policy brief. 2024. https://www.who.int/publications/i/item/9789240062764, 2026. source-page evidence item.

[9] WHO. Noncommunicable diseases. Who. noncommunicable diseases. https://www.who.int/news-room/fact-sheets/detail/noncommunicable-diseases, 2026. fact sheet. Updated April 2025.

[10] WHO Regional Office for Africa. Precision public health strategy. 2024–2030. 2024. https://www.afro.who.int/publications/precision-public-health-strategy-2024-2030, 2026. source-page evidence item.

[11] WHO. The third global patient safety challenge: tackling medication-related harm. Geneva. 2017. https://iris.who.int/bitstream/handle/10665/255263/WHO-HIS-SDS-2017.6-eng.pdf, 2026. source-page evidence item.

[12] WHO. World Malaria Report 2022. Geneva: World Health Organization. 2022. https://www.who.int/teams/global-malaria-programme/reports/world-malaria-report-2022, 2026. source-page evidence item.

Panta Rhei Research Program

Public-Good Impact Dossier

τ for Precision Public Health, Physiology-Aware Prevention, and Safer Therapeutics / Response Optimization

Dossier ID: PGID-OH-04 Portfolio: One Health 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