Contents

PROMETHEUS

A 17-Agent Mixture-of-Refinement Early-Warning System for Global Instability

---

Author: Prof. Dr. Stelios Bekiros **Contact:** bekirosstelios@gmail.com | https://www.steliosbekiros.com
Project: APOCALYPSE / PROMETHEUS **Date:** May 2026 **Status:** Public White Paper — C2.B30 Release

---

"The Index does not foretell the future. It listens to the world."

Abstract

PROMETHEUS is a neural early-warning system designed to detect the precursors of global instability before they crystallise into crises. It ingests 95 heterogeneous data channels spanning seven domains — financial and macroeconomic, social and geopolitical, solar and heliospheric, planetary and seismic, electromagnetic (Schumann resonance), climatic and disaster, and multilateral historical — and produces a single daily calibrated number, the Global Apocalypse Index (GAI), alongside a seven-day probabilistic forecast band.

The architecture is organised around 17 agents: 16 internal specialist agents, each embodying a distinct published precursor theory, and one web-facing publishing agent. The 16 internal agents are not independent models. They are doctrinal kernels within a unified Mixture-of-Refinement neural network, routed by a Bayesian mixture allocator and fused via Bayesian posterior pooling. The routing is balanced by a Sinkhorn algorithm operating under a number-theoretically derived connectivity constraint; no single precursor theory may dominate the output.

Training proceeds walk-forward over a timeline stretching from 1900 to the present, with strong regularisation and early stopping governed by a validation metric derived from temporal hold-out. A genetic algorithm outer loop, implementing multi-objective Pareto optimisation, tunes the mixture parameters after gradient training completes. The B30 canonical model runs inference on 95 channels (60 observed and 35 synthesised or frequency-aligned), with a provenance-weighted input total of 90.25 (mean fidelity 0.9124), and currently reports a GAI of 0.1348 (GREEN / STABLE).

This paper describes the system's design rationale, data architecture, agent structure, training methodology, and evaluation results at a level suitable for a scientifically literate non-specialist

audience. Technical implementation details, proprietary kernel internals, and genetic algorithm hyper-parameters are withheld; the academic publication (C7) will provide the full methodological record.

---

1. Introduction

1.1 Why Early Warning Matters

Every major global disruption — financial collapse, geophysical catastrophe, geopolitical escalation, pandemic — is preceded by observable precursors. These precursors are rarely found in a single signal. They are conjunctions: a financial stress indicator rising, a geomagnetic index becoming anomalous, social media conflict tone darkening, and a seismic cluster building — all simultaneously, in a pattern that no single-domain expert monitors end to end. Human attention is domain-local. Crises are multi-domain.

The practical consequence is that most early-warning infrastructure today operates as a collection of independent single-domain dashboards. A financial analyst watches implied volatility. A seismologist watches earthquake catalogues. A space-weather forecaster watches solar-wind data. None of these specialists has an integrated view; no existing system weighs these channels against each other in a principled, continuously updated probability estimate.

PROMETHEUS is built to address that gap directly.

1.2 Why Single-Source Signals Fail

There is a deeper theoretical reason why single-source monitoring is inadequate for catastrophic-event precursion. Catastrophic regime changes are characterised by critical slowing down: the system's internal dynamics slow, its state space contracts, and precursor signals become increasingly coherent across previously uncorrelated channels. The precursor is, in a precise mathematical sense, a cross-domain signature. A single-channel index, however well-calibrated for its own domain, cannot detect it.

Moreover, different precursor theories — Thom's catastrophe geometry, Varotsos's natural-time formalism, Hawkes's self-exciting process model, extreme value theory — make empirically distinct predictions. Their signals are not always consistent. A principled early-warning system must be able to hold all of them active simultaneously, weighting each according to its current predictive contribution, rather than committing in advance to any single theoretical framework.

1.3 Why a Mixture of Doctrinal Precursor Theories

The core design principle of PROMETHEUS is doctrinal pluralism: every established precursor theory that has empirical support in the literature is assigned a dedicated computational kernel within the network. These kernels are trained jointly, and their relative contribution to the final output is determined by the data, not by the designer. When a regime shift is approaching via a catastrophe-geometry route, the catastrophe kernel will naturally receive higher weight. When a social cascade is the primary driver, the Hawkes kernel will dominate. When multiple theories converge, the mixture allocator reflects that convergence in the output probability.

This is the Mixture-of-Refinement (MoR) principle, and it is both the technical heart and the philosophical foundation of the system.

2. Data Architecture

2.1 95-Channel B30 Panel Across 7 Domains

PROMETHEUS B30 draws on 95 data channels, deliberately distributed across seven domains so that no single domain and no single source can determine the output. The B30 build represents the fullest realisation of the "incorporate all" design principle: every key-free publicly available source was either admitted directly as observed data, or completed using statistically validated synthesis methods (described below), or excluded because it failed all completion approaches. The seven domains are:

1. **Finance and macroeconomics** — implied volatility indices, bond market stress measures, nowcast GDP estimates, economic policy uncertainty, geopolitical risk indices, global financial conditions.
2. **Social and geopolitical** — global conflict event counts and tone, social tension proxies, prediction-market sentiment (where available).
3. **Solar and heliospheric** — daily sunspot counts, planetary K-index geomagnetic measurements, solar X-ray flux, coronal mass ejection catalogues, solar-wind speed and magnetic field orientation.
4. **Planetary and seismic** — global earthquake catalogues, planetary ephemeris data, geomagnetic storm indices, early earthquake warning feeds.
5. **Schumann resonance and ELF** — measurements of the Earth-ionosphere electromagnetic resonance fundamental frequency and harmonics.
6. **Climatic and disaster** — storm event catalogues, weather anomaly feeds, international disaster databases.
7. **Multilateral and historical** — World Bank, IMF, OECD, BIS, and historical macroeconomic panels stretching back to 1900.

Of the 95 channels, 60 are directly observed through live public feeds (REAL sources). The remaining 35 are completed through a synthesis pipeline described below. All sources are publicly available or obtainable through open institutional feeds.

2.2 B30 Provenance: An Honest Accounting

The B30 design is distinguished from all prior versions by explicit, per-channel provenance tracking. Every data channel carries a label indicating how its values were obtained:

Tier	Method	What It Means
REAL	Direct observation	Values come directly from the public source
FREQ_ALIGN	Kalman, Chow-Lin, LOCF, or direct resampling	Real values at non-daily cadence aligned to the daily grid without future leakage
NOWCAST	Dynamic Factor Model or Kalman nowcast	Daily estimates inferred from co-moving real-time indicators
SYNTH	Bootstrap with matched statistics	Statistically validated synthetic series where real data is sparse

Non-observed channels enter the model's fusion layer with a weight equal to their fidelity score (mean = 0.9124 across admitted non-REAL channels). The provenance-weighted input total for the B30 panel is 90.25 out of a possible 95.

Every non-REAL channel must pass a battery of 40 statistical tests covering moments, autocorrelation structure, tail properties, and cross-correlations with related real channels before it is admitted. Sources that fail all synthesis methods are excluded and documented in an honest list of 35 truly-failed sources.

2.3 The Polarity-Validated Panel

A critical pre-training verification step is the polarity diagnostic: every source in the panel must correlate positively with observed crisis severity in the historical record. After a direction-tagging pass — which inverts sources where raw values decrease during crises (for example, the Dst geomagnetic storm index and the GDELT global tone score) — every production source passes the polarity check before being admitted. This ensures that every channel points in the same semantic direction: higher value means higher risk.

Sources that fail network access during any given inference run are masked rather than imputed; the network is trained to operate on partially available panels.

2.4 Rolling Five-Year Normalisation

Under previous data layers, per-source normalisation was computed from each training window in isolation. On a calm window the recent low variance dominated the scale, mapping crisis-grade values toward the mid-band and silently degrading signal. The current data layer (introduced at B27) replaces this with rolling five-year cached median/MAD per source, anchor-aware for sources that have analytic anchors, so short-window inference and long-window training share one risk scale. A per-cell freshness/age channel records staleness in days normalised by the source's expected cadence, and publication-lag enforcement guarantees that the value visible at date d is the value known by date d, eliminating look-ahead leakage.

2.5 B30 Panel Composition Summary

- Total admitted channels: **95**
- REAL channels: **60** (including real backfill)
- FREQ_ALIGN channels: **14** (8 Kalman + 4 LOCF + 1 Chow-Lin + 1 Direct)
- NOWCAST channels: **15** (14 DFM + 1 Kalman)
- SYNTH channels: **6** (bootstrap)
- Excluded (truly failed): **35** sources
- Provenance-weighted total: **90.25**
- Mean fidelity score: **0.9124**

---

3. Agent Architecture

PROMETHEUS comprises 17 agents: 16 internal agents that embody distinct precursor theories, and one outward-facing agent responsible for publishing.

The 16 internal agents are not 16 separate neural networks. They are 16 specialised computational kernels within a single unified model, each receiving the same time-series panel, each applying a distinct mathematical lens, and each contributing a posterior probability estimate to a shared mixture output.

Agent 01 — Super Intuition (Orchestrator)

This is the fusion stage. Agent 01 holds the Bayesian posterior fusion layer that pools the 16 expert outputs into a single calibrated probability. Each expert's log-odds estimate is weighted by a learned set of coefficients constrained to the probability simplex, then combined with a Dirichlet prior that prevents any single expert from achieving total dominance.

Agent 02 — Catastrophe (Cusp Geometry)

Grounded in Rene Thom's catastrophe theory, this kernel detects approach to fold and cusp bifurcations in the system's latent dynamics. Catastrophe theory provides a geometric framework for sudden qualitative regime changes — the cusp catastrophe captures the hysteresis typical of financial crashes and geopolitical tipping points. When this kernel's contribution weight rises sharply, it signals that the system's geometry is approaching a boundary beyond which small perturbations produce large, discontinuous responses.

Agent 03 — Chaos (Lyapunov / Fractal Dimension)

This kernel quantifies the sensitivity of the system's current state to perturbations — the foundational concern of chaos theory. It computes a differentiable approximation of the fractal dimension of the current trajectory and a finite-time Lyapunov exponent estimate over a sliding window. As a complex system approaches a critical transition, its attractor geometry typically contracts and its Lyapunov spectrum changes.

Agent 04 — Recurrent (Recurrence Quantification Analysis)

Grounded in the recurrence quantification analysis framework, this kernel measures how often the system revisits states it has previously occupied. The key statistics are determinism (the fraction of recurrence points forming diagonal lines) and laminarity (the fraction forming vertical lines). These statistics change characteristically before critical transitions.

Agent 05 — Genetic (Evolutionary Optimiser)

This agent operates at training time, not at inference time. It implements the multi-objective genetic algorithm outer loop — NSGA-II non-dominated sorting, ALPS age-layered population structure, BLX-alpha arithmetic crossover, and Cauchy heavy-tailed mutation — that optimises the mixture allocation parameters and zone classification thresholds after gradient training. Agent 05 is the mechanism by which the system discovers the Pareto-optimal trade-off between detection rate, false-alarm rate, and calibration, without any human-set weights.

Agent 06 — Extreme Value Theory (Tail Risk)

Grounded in the Fréchet-Tippett-Gnedenko extreme value theorem and the generalised extreme value distribution family, this kernel estimates the probability that an observed value exceeds any given threshold. Its role is specifically to detect situations where the deterministic precursor kernels report relative calm but the statistical tail of the underlying distribution is materially elevated. It provides formal early warning of fat-tail events.

Agent 07 — Indicators (Technical and Statistical Feature Bank)

This is the broadest kernel, computing a wide bank of rolling statistical features from the full panel: return series, realised volatility, drawdowns, cross-source z-scores, term-structure spreads, and momentum-like acceleration measures. It is the baseline against which the more theory-specific kernels add their corrections and is almost always active.

Agent 08 — Helix (Rotational Residual / Cyclic Structure)

This kernel captures coupled rotation and translation in the latent trajectory — a structural signature of long-memory cyclic processes including planetary orbital configurations, solar cycles, and social rhythms. The mathematical construction uses an SO(2) rotation applied to the residual stream at a learned phase angle, providing a phase-aware path that ordinary skip connections cannot represent.

Agent 09 — Natural Time (Criticality via Varotsos-Sarlis-Skordas)

The natural-time formalism provides an alternative to clock time for the analysis of event sequences. In natural time, the order of events matters but their clock-time spacing does not. The key statistic, kappa-1, is the variance of the natural-time index, and it approaches a universal critical value near 0.07 as a system enters the critical state preceding a seismic, financial, or other collective event.

Agent 10 — Ramanujan (Number-Theoretic Activation)

This kernel applies the Rogers-Ramanujan continued fraction expansion as a modulated activation function within the refinement blocks, and computes a structural roughness signature on the incoming time series using number-theoretic correlation functions. The Rogers-Ramanujan product formula generates a smooth, saturating non-linearity with distinctive curvature, empirically shown to reduce numerical instability in deep configurations.

Agent 11 — Schumann (Earth-Ionosphere Electromagnetic Resonance)

The Schumann resonances are the set of electromagnetic resonances of the Earth-ionosphere cavity, with a fundamental frequency near 7.83 Hz. Their amplitude, frequency, and quality factor respond to global lightning activity and large-scale geomagnetic disturbances. This kernel applies a multi-level wavelet decomposition to multi-station Schumann measurements and extracts the cross-station coherence structure.

Agent 12 — Social (Hawkes Self-Exciting Process)

Grounded in the Hawkes self-exciting point process model, this kernel treats conflict events and social tension escalations as sequences with self-excitation: each event increases the probability of further events in the near future, with a decaying memory kernel. The GDELT global conflict event database provides the primary input where available.

Agent 13 — Solar (Heliospheric Precursors)

This kernel synthesises multiple solar and space-weather data streams: daily sunspot counts, planetary K-index geomagnetic activity, solar X-ray flux, coronal mass ejection catalogues, and solar-wind speed and magnetic field measurements. A transit-time alignment step shifts the solar-wind measurements by the estimated travel time from the Sun to Earth before merging them with terrestrial observations.

Agent 14 — Planetary (Seismic and Orbital Signals)

This kernel combines seismology — applying the natural-time criticality measure to decluttered global earthquake catalogues — and orbital mechanics, computing synodic angular separation sequences for key planetary pairs and applying the natural-time framework to those sequences. The Dst geomagnetic storm index is also ingested as a proxy for magnetospheric response to solar-wind forcing.

Agent 15 — Wavelets (Multi-Resolution Decomposition)

Grounded in wavelet theory, this kernel applies a stationary discrete wavelet transform to the panel, selecting the decomposition depth that minimises Shannon entropy and the wavelet basis that achieves the sparsest representation. Multi-scale residuals are fed both to the shared feature stem and to the wavelet expert head itself.

Agent 16 — Data (I/O Substrate)

Agent 16 does not contribute a precursor kernel. It is the input-output substrate: it owns the source registry, the adapter framework for fetching and normalising each of the 95 channels in the B30 panel, and the masking logic that handles missing or unavailable data.

Agent 17 — Website (Public Interface)

The 17th agent is the public-facing portal. It is not part of the inference graph; it is the sole consumer of the model's output. Every day at 12:00 GMT, it reads the latest GAI value, zone classification, seven-day forecast band, and top contributing kernels, and publishes them to the public dashboard at `https://apocalypse-bekiros.pplx.app`.

---

4. The Mixture-of-Refinement Core

The 16 doctrinal kernels are the experts in a Mixture-of-Refinement architecture. The key design challenge for any mixture model is preventing expert collapse: the tendency for routing mechanisms to converge on one or two dominant experts and ignore the rest. PROMETHEUS addresses this through two structural innovations.

The first is the routing mechanism itself. Rather than a simple top-k hard assignment, the system uses a Bayesian Latent Mixture Allocator that performs soft assignment across all 16 experts, balanced by a Sinkhorn iteration that enforces approximate load balancing. No expert is ever completely silenced. No expert can completely monopolise the output: a dominance guard imposes a ceiling on the fraction of the posterior any single kernel may contribute.

The second is the attention connectivity structure between experts. The inter-agent information exchange is governed by a connectivity pattern derived from Ramanujan expander graph theory, which provides a provably near-optimal balance between information propagation speed and propagation balance. This prevents the routing mechanism from developing structural preferences for particular experts independent of the data.

The fusion stage combines the routed expert outputs in log-odds space with a learned Dirichlet prior, producing a single calibrated probability that is directly interpretable as a statement about global instability risk.

---

5. Training and Evaluation Methodology

5.1 Walk-Forward Training

PROMETHEUS is trained in walk-forward mode over a timeline from 1900 to the present. Walk-forward training means that at each training step, the model sees only data from the past relative to the forecast point — never information from the future. The validation set is the most recent period of the timeline, held out entirely from gradient training. The primary evaluation metric — AUC — is computed on this temporal hold-out and tracked with an exponential moving average to smooth epoch-to-epoch fluctuation.

5.2 Stratified Validation

A critical improvement introduced at B27 is the stratified validation split. The earlier chronological 90/10 split reserved the final ten percent of the timeline for validation; on calm regimes this tail contained zero positive (crisis) labels, which collapses AUC and ECE to noise. The current system uses a stratified random sample with both crisis and non-crisis periods represented in validation, giving the evaluation metrics their proper diagnostic power.

5.3 Regularisation and Early Stopping

The canonical training configuration employs strong weight decay (5e-2, which is 50x heavier than the early B-track default) to resist memorisation of historical crisis patterns. Early stopping is triggered when the held-out performance fails to improve for three consecutive epochs. A composite loss function combines a Brier probability score, a focal loss component, a lead-time reward, and a Bayesian KL-regularisation term.

5.4 Genetic Algorithm Tuning

After gradient training, a genetic algorithm outer loop performs multi-objective Pareto optimisation over the mixture allocation parameters and zone classification thresholds. The three objectives being simultaneously optimised are: maximising the area under the ROC curve, maximising the detection lead time, and minimising the false-alarm rate. No human-set weights govern this trade-off.

5.5 Two-Stage Calibration

Post-training calibration proceeds in two stages. The first is Platt sigmoid scaling with a bias grid search over 80 x 25 = 2,000 parameter combinations, selecting the pair that minimises held-out expected calibration error. The second is pool-adjacent-violators isotonic regression on the raw probabilities, which enforces a monotonicity constraint. The version achieving lower held-out ECE is deployed.

5.6 Three-Meter Sign-Off Gate

The system is not deployed until it passes all three of the following metrics:

1. **AUC on temporal hold-out** – measures the system's ability to rank crisis periods above non-crisis periods. Required: >= 0.70.
2. **False-alarm rate at 80% detection coverage (FAR@POD=0.80)** – measures how many non-crisis periods are incorrectly flagged when the system is operating at 80% detection sensitivity. Required: <= 0.20.
3. **Expected calibration error (ECE)** – measures how closely the system's stated probability estimates correspond to empirical frequencies. Required: <= 0.10.

---

6. B30 Results

6.1 B30 Build Overview

The B30 build was completed on 2026-05-06. Its distinguishing feature is the 95-channel provenance-aware panel, which replaces the S=47 panel of B23 and the S=60 real panel of B27. The B30 panel introduces 35 synthesised or frequency-aligned channels, each carrying explicit provenance metadata and each validated through a 40-test statistical fidelity battery.

6.2 Canonical B30 GAI

At the time of writing, the B30 canonical checkpoint (MD5 `3a4244b32f1976dc21eb33bb7d3897d4`) produces a GAI of **0.1348** in the **GREEN / STABLE** zone. The seven-day peak forecast is 0.2081, indicating no near-term elevation above the green threshold on current trajectories.

6.3 Backtest AUC

Backtest average AUC across walk-forward windows on the current B30 panel: **0.833**. This reflects genuine out-of-sample prospective performance under strict temporal causality at every step.

6.4 Provenance Composition

The B30 panel has a provenance-weighted input total of **90.25** out of a possible 95. The gap (4.75) reflects the 35 truly-failed sources that are excluded, and the downweighting (below 1.0) applied to non-REAL channels proportional to their fidelity score. Mean fidelity across non-REAL admitted channels: **0.9124**.

The provenance composition board on the public dashboard provides a real-time breakdown by tier.

6.5 A Note on B30 Promotion

The B30 retrain generated nine candidate checkpoints. Candidate H was promoted under a documented one-time Path 3 exception applied specifically to the panel-expansion shock (47 -> 95 channels), where AUC and ECE passed and FAR was within a pre-documented recalibration band. **The 0.20 FAR threshold is locked for all future daily retrains.** This transitional exception applied once, for the B30 build only; all subsequent retrains must pass the full three-meter gate at its documented thresholds.

6.6 Previous Canonical Results (B27, for Reference)

The B27 canonical checkpoint, promoted 2026-05-05, achieved:

Meter	Value	Threshold	Pass
AUC_val(EMA)	0.906	>= 0.70	Yes
FAR@POD=0.80	0.032	<= 0.20	Yes
ECE	0.096	<= 0.10	Yes

---

7. Limitations

Transparency about limitations is as much a part of the system's design philosophy as the methodology itself.

Train-validation gap. The difference between performance on training data and on held-out validation data remains larger than ideal. Even with strong regularisation, the system shows signs of overfitting to the distribution of historical crisis patterns. Architectural dropout — not yet deployed — is planned to address this.

Coverage gaps in the B30 panel. Thirty-five sources failed all completion methods and are excluded from the B30 panel. Notably, all five Schumann resonance sources and the primary social conflict source (GDELT) are absent from the B30 run due to endpoint failures on the build date. Prediction-market sources (Kalshi, Polymarket, Manifold, Metaculus) also failed. The affected kernels continue to operate on historical training signal but do not receive live data for those domains in B30.

Synthesised channels are not observations. Thirty-five of the 95 B30 channels are synthesised or aligned rather than directly observed. Each is validated through the 40-test fidelity battery and carries a fidelity weight in the fusion layer, but they are not interchangeable with real observations. The provenance composition board on the dashboard makes this visible. As live coverage of these sources improves, they will be promoted to REAL status.

This is not a prediction of specific events. The Global Apocalypse Index is a probabilistic instability index, not a forecast of any particular event. It estimates the aggregate probability that the global system is in an elevated instability regime. The distinction is fundamental: the system is a fire-danger index, not a fire-location predictor.

Historical calibration is not future performance. The system's performance metrics are computed over historical data. The nature, frequency, and inter-domain correlation structure of future crises may differ from historical precedent. The walk-forward training protocol is designed to minimise train-test leakage, but no historical evaluation can fully anticipate novel crisis configurations.

The model is updated daily. Between daily retraining runs, the model's parameters are fixed. This means that rapidly evolving situations that have no historical analogue may be underweighted in the immediate post-event window, until the next daily retrain incorporates them.

---

8. Use Cases

8.1 Public Awareness and Dashboards

The primary intended use of the GAI is as a public information instrument — a continuously updated, machine-generated assessment of global instability conditions, presented with full transparency about its inputs, methodology, and limitations. The META-PORTAL (C3) is the public face of this function, updated daily at 12:00 GMT.

8.2 Policy Briefings and Situational Awareness

Analysts, think tanks, and policy institutions can use the GAI and its domain sub-indices to characterise the current instability environment in structured quantitative terms. The system's attribution output — which kernels are contributing most to the current reading — provides a domain-level decomposition that supports briefing narratives.

8.3 Academic and Scientific Discourse

PROMETHEUS is designed to be a contribution to the empirical literature on multi-domain precursor detection. The academic publication (C7) will provide the full methodological specification for replication and critique. The system's architecture instantiates testable hypotheses about which precursor theories carry signal in which regimes.

8.4 Financial and Actuarial Overlays

The GAI and its constituent domain indices can serve as conditioning variables in financial stress analysis and actuarial risk modelling. No specific financial product or trading strategy is implied or endorsed.

---

9. Roadmap

Architectural dropout. Adding stochastic regularisation within each kernel's internal layers is expected to meaningfully reduce the train-validation gap and improve the robustness of probability estimates.

Tier 2 capacity expansion. A larger model configuration — approximately 100 million parameters, compared to the current 38 million — is gated on further validated performance improvements from the dropout pass.

Live source recovery. The five Schumann sources, GDELT, and the prediction-market feeds that failed on the B30 build date are targets for re-admission. As each recovers or an alternate source is wired, it will be promoted from SYNTH to FREQ_ALIGN or REAL status.

Keyed source binding. Provisioning API credentials for FRED, prediction-market feeds (Kalshi, Polymarket, Manifold, Metaculus), and major social-media data streams will bring live adapter coverage toward 70 of 95 channels.

Conformal prediction coverage. The seven-day forecast band will be wrapped with conformal prediction guarantees, providing distribution-free coverage bounds.

Model ensembling. Once the Tier 2 model is trained and validated, calibration-preserving ensembling of the Tier 1 and Tier 2 checkpoints will be implemented.

Migration to steliosbekiros.com. When migration to Hetzner is approved, the dashboard will be served at `apocalypse.steliosbekiros.com` from the same JSON feeds, with the daily heartbeat running as systemd timers on a self-hosted box.

---

10. References

Catastrophe Theory

Thom, R. (1975). *Structural Stability and Morphogenesis: An Outline of a General Theory of Models*. W. A. Benjamin.

Zeeman, E. C. (1977). *Catastrophe Theory: Selected Papers, 1972-1977*. Addison-Wesley.

Natural Time Analysis

Varotsos, P. A., Sarlis, N. V., & Skordas, E. S. (2011). *Natural Time Analysis: The New View of Time*. Springer.

Sarlis, N. V., Skordas, E. S., & Varotsos, P. A. (2010). Nonextensivity and natural time: The case of seismicity. *Physical Review E*, 82(2), 021110.

Hawkes Processes and Self-Exciting Models

Hawkes, A. G. (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1), 83-90.

Ogata, Y. (1988). Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association*, 83(401), 9-27.

Recurrence Quantification Analysis

Marwan, N., Romano, M. C., Thiel, M., & Kurths, J. (2007). Recurrence plots for the analysis of complex systems. *Physics Reports*, 438(5-6), 237-329.

Eckmann, J.-P., Kamphorst, S. O., & Ruelle, D. (1987). Recurrence plots of dynamical systems. *Europhysics Letters*, 4(9), 973-977.

Extreme Value Theory

Embrechts, P., Kluppelberg, C., & Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance*. Springer.

Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer.

Antal, A., Droz, M., Gyorgyi, G., & Racz, Z. (2001). 1/f Noise and extreme value statistics. *Physical Review Letters*, 87(24), 240601.

Mixed-Frequency Methods and Temporal Disaggregation

Chow, G. C., & Lin, A. L. (1971). Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *Review of Economics and Statistics*, 53(4), 372-375.

Ghysels, E., Santa-Clara, P., & Valkanov, R. (2004). The MIDAS touch: Mixed data sampling regression models. *CIRANO Working Papers*.

Harvey, A. C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

Wavelet Methods

Mallat, S. (2009). *A Wavelet Tour of Signal Processing: The Sparse Way* (3rd ed.). Academic Press.

Daubechies, I. (1992). *Ten Lectures on Wavelets*. SIAM.

Schumann Resonance

Nickolaenko, A., & Hayakawa, M. (2002). *Resonances in the Earth-Ionosphere Cavity*. Springer.

Price, C. (2016). ELF electromagnetic waves from lightning: The Schumann resonances. *Atmosphere*, 7(9), 116.

Ramanujan and Number-Theoretic Methods

Hardy, G. H., & Ramanujan, S. (1918). Asymptotic formulae in combinatory analysis. *Proceedings of the London Mathematical Society*, Series 2, 17, 75-115.

Vaidyanathan, P. P., & Tenneti, S. (2014). Ramanujan filter banks for estimation and tracking of periodicities. *IEEE ICASSP 2014 Proceedings*, 3851-3855.

Critical Slowing Down and Complex Systems

Scheffer, M., Bascompte, J., Brock, W. A., et al. (2009). Early-warning signals for critical transitions. *Nature*, 461(7260), 53-59.

Dakos, V., Scheffer, M., van Nes, E. H., et al. (2008). Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences*, 105(38), 14308-14312.

Mixture-of-Experts and Neural Routing

Jacobs, R. A., Jordan, M. I., Nowlan, S. J., & Hinton, G. E. (1991). Adaptive mixtures of local experts. *Neural Computation*, 3(1), 79-87.

Shazeer, N., Mirhoseini, A., Maziarz, K., et al. (2017). Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv:1701.06538*.

Geomagnetic and Planetary Data

Sugiura, M. (1964). Hourly values of equatorial Dst for the IGY. *Annals of the International Geophysical Year*, 35, 9-45.

Bayesian Calibration

Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 359-378.

Niculescu-Mizil, A., & Caruana, R. (2005). Predicting good probabilities with supervised learning. *Proceedings of ICML 2005*, 625-632.

---

11. Author Biography

Prof. Dr. Stelios Bekiros is Chair Professor of Finance, Economics and Artificial Intelligence, with research affiliations across European and North American institutions, including positions in econometrics, nonlinear dynamics, and machine learning for complex financial and physical systems. His work spans catastrophe theory, multi-agent computational intelligence, extreme value analysis, and the application of deep learning to non-stationary multi-source time series. APOCALYPSE / PROMETHEUS is the synthesis of his research programme in these areas.

Contact: bekirosstelios@gmail.com | https://www.steliosbekiros.com

---