FICSS WORKING PAPER | WP-2026-005

Beyond the Lyapunov Wall:

Extending the Predictive Horizon of Computational Macrohistory

A Methodological Roadmap

Galen Fontaise

Fontaise Institute of Computational Social Science (FICSS), Lugano, Switzerland

April 2026

Abstract

The Computational Macrohistory (CMH) framework, as developed in the preceding papers of this series, operates within a temporal horizon of roughly five to fifteen years for quantitative probabilistic forecasts. This constraint is not a technological artefact to be overcome with faster computation or larger datasets: it is a structural consequence of chaotic dynamics formalised in Axioms A5 and A7 of the CMH axiomatic system, and known in the broader literature on dynamical systems as the Lyapunov Wall. The present paper asks whether, and to what extent, mathematical strategies exist that can push this horizon further without abandoning the epistemological rigour on which CMH is founded.

A central claim of this paper is that the question of horizon extension cannot be answered without first specifying what kind of prediction is being considered. We distinguish three levels of the CMH prediction space: event prediction (the trajectory-level forecast of system state at a future time, subject to the full force of the Lyapunov Wall); regime probability (the probability that the system will be operating in a given dynamical regime over a future interval, subject to a substantially slower decay); and spectral structure persistence (the identification of invariant or near-invariant cyclic structures in the system's dynamics, whose predictive content is not subject to exponential decay but to linear phase uncertainty and structural stability constraints). The five methodological strategies surveyed in this paper operate at different levels of this hierarchy, and the comparison is meaningful only when the level is held constant.

We survey five candidate approaches — Ensemble Methods, Slow Manifold theory combined with Critical Slowing Down indicators, Analog Forecasting, Reservoir Computing, and Koopman Operator theory implemented through Dynamic Mode Decomposition (DMD) — organising them into three methodological classes: uncertainty propagation strategies, regime detection strategies, and spectral decomposition strategies. We identify the Koopman Operator / DMD framework as the most promising candidate for integration with CMH at the level of spectral structure persistence, owing to its natural connection with the ergodicity assumption formalised in CMH Axiom A2, its compatibility with the panel structure of historical datasets, and its capacity to produce empirically falsifiable predictions about persistent cyclic structures.

The paper concludes with a description of the historical datasets most suited to a first empirical test of Koopman-CMH, including explicit criteria for dataset selection, and an announcement of the companion

paper that will develop the full mathematical treatment.

Keywords: *Computational Macrohistory, Lyapunov Wall, Koopman operator, Dynamic Mode Decomposition, long-range forecasting, complex systems, historical dynamics*

JEL Classification: C02, C63, N00, D74

FICSS Working Paper Series No. WP-2026-005. Fontaise, G. (2026). Beyond the Lyapunov Wall: Extending the Predictive Horizon of Computational Macrohistory. A Methodological Roadmap. FICSS Working Paper WP-2026-005. Fontaise Institute of Computational Social Science, Lugano.

1. INTRODUCTION

Every science of forecasting confronts a temporal horizon beyond which its predictions become unreliable. In meteorology the horizon is approximately two weeks. In seismology the precise timing of earthquakes remains essentially unknowable. In economics, consensus macroeconomic forecasts lose most of their marginal value beyond two to three years. These limits are not expressions of ignorance to be dispelled by accumulating more data or building faster computers. They are structural properties of the systems being studied, rooted in the mathematics of non-linear dynamics.

The Computational Macrohistory (CMH) framework, developed systematically in WP-2026-001 through WP-2026-004 of this series, is no exception. Its axiomatic foundations explicitly acknowledge that predictive accuracy decays monotonically with the temporal horizon (Axiom A7) and that this decay is driven by the endogenous indeterminacy of chaotic dynamics (Axiom A5). The practical consequence is that CMH's quantitative probabilistic forecasts are reliable for horizons of five to fifteen years — a range sufficient for many policy-relevant applications, but insufficient for the broader ambition of a science of large-scale historical dynamics.

The question this paper addresses is therefore not whether the Lyapunov Wall can be eliminated — it cannot — but whether it can be partially circumvented by changing the nature of the predictive question. This reframing is not a rhetorical convenience. It reflects a substantive insight: the Wall is a constraint on a specific kind of prediction, namely the trajectory-level forecast of where a system will be at a future point in time. It is a much weaker constraint — and in some formulations no constraint at all — on predictions about the statistical structure of the system's long-run behaviour, the probability of being in one regime rather than another, or the persistence of cyclic patterns in the dynamics.

To make this precise, we introduce a three-level hierarchy of the CMH prediction space that will organise the entire comparative analysis. At the first level sits event prediction: the trajectory-level forecast that asks where exactly the system will be at time $t + \tau$. This is the prediction type that A7 governs directly, and its accuracy decays exponentially with τ at a rate bounded below by the maximum Lyapunov exponent. At the second level sits regime probability: the forecast of the probability that the system will be operating in a qualitatively defined dynamical regime — stable, unstable, approaching transition — over a future interval. This type decays more slowly, because it requires only that the system's position in state space be approximately correct at a coarse level of resolution. At the third level sits spectral structure persistence: the identification of invariant or near-invariant cyclic structures in the dynamics, the confirmation that such cycles exist, and the estimation of their periods. For truly invariant structures, no decay applies, because the question is not about the system's location but about the geometry of the space it inhabits.

The five methodological families surveyed in this paper operate at different levels of this hierarchy.

Comparing them as if they were alternative answers to the same question would be misleading; the comparison is meaningful only when the prediction level is held explicit. The paper proceeds as follows. Section 2 provides an accessible account of the Lyapunov Wall and its CMH-specific quantification. Section 3 surveys the five candidate approaches, organised into three methodological classes, and evaluates each against the prediction hierarchy. Section 4 presents the Koopman Operator framework as the leading candidate at the level of spectral structure persistence. Section 5 describes the dataset selection criteria and the historical sources most suited to an initial empirical test. Section 6 discusses expected contributions and limitations. Section 7 concludes.

2. THE LYAPUNOV WALL: A STRUCTURAL LIMIT, NOT A TECHNOLOGICAL ONE

The term "Lyapunov Wall" refers to the finite predictive horizon that emerges in any non-linear dynamical system exhibiting sensitive dependence on initial conditions. The concept is named after the Russian mathematician Aleksandr Lyapunov, whose late nineteenth-century work on the stability of differential equations laid the foundation for what later became chaos theory. The key quantity — the maximum Lyapunov exponent, typically denoted λ_{max} — measures how rapidly two trajectories that begin arbitrarily close together diverge over time.

Consider two weather systems beginning in states that differ by an imperceptible amount — say, a marginal difference in atmospheric pressure over the Atlantic. In a stable system, this difference would remain small or shrink to zero. In a chaotic system, however, the difference grows exponentially: after a characteristic time proportional to the inverse of λ_{max}, the two trajectories have diverged so far that they bear no resemblance to one another. Any forecast based on the initial state has therefore become meaningless. For Earth's atmosphere, this characteristic time is approximately two weeks — which is why meteorological forecasts beyond that horizon have essentially no skill.

For socio-political systems of the kind studied by CMH, direct empirical estimation of maximum Lyapunov exponents at the macro-political level remains methodologically challenging and the existing literature is sparse. Studies of socio-economic time series — including analyses of economic fluctuation dynamics (Medio and Gallo, 1995) and of financial market behaviour (Gao et al., 2007) — report positive Lyapunov exponents broadly consistent with a range of 0.2 to 0.5 per year, though these estimates derive from systems that differ in important respects from the macro-political dynamics that CMH targets. In the absence of more direct measurements for political instability dynamics specifically, this range should be understood as an order-of-magnitude working hypothesis rather than an empirically settled value; developing a CMH-specific Lyapunov estimation methodology is an open research task identified for the companion paper. With this qualification, the values can be given a precise operational meaning within the CMH framework. The useful predictive horizon τ^* — the time at which forecast uncertainty grows from the measurement error on the CMH state variables to the macro-scale threshold that distinguishes one regime from another — satisfies:

$$\tau^* \approx (1 / \lambda_{max}) \cdot \log(\Delta / \delta_0)$$

where δ_0 is the observational uncertainty in the initial state of the CMH variables and Δ is the macro-scale resolution relevant for the prediction. Both quantities can be given concrete operational definitions within the CMH framework. δ_0 is the typical measurement discrepancy for each state variable across independent data sources: for the Polity democracy score, cross-source discrepancies between Polity V and V-Dem for the same country-year typically fall in the range of ±1 to 2 points on a 20-point scale, corresponding to a relative uncertainty of approximately 5–10%. For the Gini coefficient, inter-source discrepancies are of similar order. Δ is the threshold in the Systemic Stress Index that separates stable from unstable regime classifications, as operationalised in WP-2026-004: the SSI threshold of 0.5 on the normalised scale corresponds to a macro-scale resolution that is approximately one to two orders of magnitude larger than δ_0, placing δ_0 / Δ in the range 10^{-2} to 10^{-3}. Substituting these values and the working hypothesis $\lambda\_max \in [0.2, 0.5]$ per year yields τ^* in the range of five to fifteen years for event-level prediction — a range that is now operationally grounded rather than merely asserted, even if its precise boundaries remain subject to empirical revision.

This limitation is irreducible at the level of event prediction — a point that is frequently misunderstood. Collecting more data does not help, because the problem lies not in insufficient knowledge of the initial state but in the mathematical structure of the system's evolution. More powerful computers do not help either, because even a perfect simulation of a chaotic system eventually loses track of the true trajectory. And better models do not dissolve the Wall, because the exponential divergence is a property of the actual system, not of any particular approximation to it.

What can be done is to ascend the prediction hierarchy introduced in Section 1. Regime-level and spectral-level predictions are subject to different, and generally much slower, forms of decay. The five approaches surveyed in Section 3 each implement a different strategy for this ascent. Before proceeding, however, it is worth emphasising a risk that this paper must guard against: the appearance of circumventing A7 by relabelling description as prediction. We therefore maintain throughout a strict distinction between identifying a structural pattern in historical data — which is description — and using that pattern to generate a falsifiable probabilistic statement about future behaviour — which is prediction. Only the latter counts as a contribution to the CMH predictive programme.

3. CANDIDATE APPROACHES: A COMPARATIVE ASSESSMENT

The five methodological families considered here are not homogeneous in their theoretical foundations: one is a strategy for propagating uncertainty through ensemble simulation, two are methods for detecting regime structure and impending transitions, and two are spectral or machine-learning approaches that seek persistent representations of the dynamics. Treating them as competing answers to the same question would obscure more than it illuminates. We therefore organise the comparison into three classes, following the prediction hierarchy of Section 1, and assess each approach on the four criteria of temporal horizon, data requirements, applicability to the CMH state-space architecture, and methodological maturity. Table 1 summarises the results.

3.1 Class I — Uncertainty Propagation: Ensemble Methods and Analog Forecasting

Class I approaches do not change the level of prediction: they remain at the event-prediction level, but characterise the uncertainty around that prediction more faithfully. Their contribution is a better-described Wall, not a higher one. Both methods in this class are methodologically mature and widely deployed; their inclusion in this comparative assessment is warranted not because they extend the predictive horizon in any fundamental sense, but because they constitute necessary infrastructure for any rigorous CMH forecasting system and provide a quantitative baseline against which the more ambitious Class III approaches must be measured.

Ensemble methods are the most widely deployed strategy for managing forecast uncertainty in chaotic systems. Rather than producing a single trajectory from a single set of initial conditions, an ensemble model generates many trajectories — typically dozens to hundreds — each starting from a slightly different initial state drawn from a probability distribution that reflects the observational uncertainty in the initial conditions. The spread of the resulting ensemble provides a direct, empirical measure of forecast uncertainty: a tightly clustered ensemble indicates a configuration in which small differences in initial conditions produce similar futures — a predictable situation; a rapidly diverging ensemble signals that small initial differences produce qualitatively different futures — intrinsic unpredictability regardless of model quality. The European Centre for Medium-Range Weather Forecasts, which runs the world's most capable operational ensemble prediction system, uses fifty-one members and provides calibrated probabilistic forecasts up to two weeks with demonstrated skill. The gain over the 1970s, when reliable forecasts extended to only three days, is largely attributable to ensemble methods combined with improvements in observation density — not to any fundamental advance in the physics of the atmosphere or in the mathematics of prediction.

For Computational Macrohistory, ensemble simulation is already implicit in the Monte Carlo framework described in WP-2026-002: multiple trajectory integrations of the CMH state equation, each with perturbed initial conditions drawn from the measurement uncertainty distributions of the state variables, yield an empirical distribution over future states that is the natural output of the CMH predictive system. What ensemble methods add is a principled calibration of this distribution — ensuring that the spread of the ensemble corresponds to the actual frequency of outcomes, a property known as reliability or calibration. A well-calibrated ensemble that says a thirty percent probability of severe instability should be wrong about seventy percent of the time in those cases. Achieving this requires careful tuning of the perturbation scheme and ongoing validation against held-out data. For CMH applications requiring horizons of decades or centuries, ensemble methods remain necessary but not sufficient: they characterise uncertainty up to the Wall, but they cannot push the Wall back.

Analog forecasting, proposed by Lorenz himself in 1969, takes a fundamentally different approach. Rather than simulating future trajectories from the current state, it searches the historical record for past states of the system that closely resemble the current state — the analogs — and uses the subsequent evolution of those past states as a forecast for the current trajectory. The appeal of the method is considerable: it requires no explicit model, makes no assumptions about the functional form of the dynamics, and automatically captures non-linearities that a parametric model might miss. In practice, the method works well in domains where the historical archive is large enough to contain close analogs — atmospheric reanalysis datasets, which now span decades at high temporal and spatial resolution, provide sufficient density for analog forecasting to compete with numerical weather prediction on short horizons.

For CMH, the fundamental barrier to analog forecasting is the curse of dimensionality. The volume of state space grows exponentially with the number of variables: a system with twenty-five variables, each requiring at minimum ten distinct values to characterise its range adequately, would require ten to the power of twenty-five distinct historical observations to populate the state space at even the coarsest resolution. The historical record of macro-political variables, even at its most ambitious, contains on the order of tens of thousands of country-year observations — a number that falls orders of magnitude short of what is needed for the twenty-five-dimensional CMH space. Dimensionality reduction through principal component analysis or similar techniques can partially address this problem by working in a lower-dimensional latent space rather than the full state space, but at the cost of discarding the precise combinations of variables that may be most predictive of specific transition types. Analog forecasting thus remains a conceptually valuable benchmark method for CMH — useful for identifying historical cases that structurally resemble a target case — but not a viable path to extended predictive horizons at the level of ambition that motivates this paper.

3.2 Class II — Regime Detection: Slow Manifold Theory and Critical Slowing Down

Class II approaches operate at the second level of the prediction hierarchy. They do not predict where the system will be, but rather what regime it will be in and — most usefully — whether it is approaching a qualitative transition. The decay of this type of prediction is governed not by the full Lyapunov exponent but by the much slower dynamics of the system's structural variables. This distinction between fast and slow variables, which is implicit in the CMH state-space architecture, becomes the explicit mathematical foundation of both methods in this class.

Slow manifold theory, developed in its modern rigorous form by Fenichel (1979) and grounded in the earlier work of Tikhonov (1952) on singularly perturbed differential equations, formalises the observation that many complex systems contain variables evolving on radically different time scales. When the ratio of fast to slow time scales is sufficiently large — a condition known as the singular perturbation assumption — the fast variables reach a quasi-equilibrium almost instantaneously relative to the slow variables, and the relevant long-run dynamics are entirely captured by the evolution of the slow subsystem on a lower-dimensional manifold embedded in the full state space. The key result, Fenichel's theorem, guarantees that this slow manifold exists, is smooth, and persists under small perturbations, provided the time scale separation is sufficiently pronounced.

The CMH state-space architecture maps naturally onto this framework. The demographic variables — population growth rate, youth bulge, urbanisation rate, dependency ratio — evolve on characteristic time scales of fifteen to thirty years, determined by biological constraints on reproduction, ageing, and migration. The economic structural variables — GDP per capita trend, Gini coefficient, fiscal capacity — evolve on scales of five to fifteen years, driven by capital accumulation, institutional change, and distributional politics. The political event variables — Polity score fluctuations, government longevity — have both slow structural components and fast event-driven components, the latter operating on scales of months to a few years. The collective psychology variables are the fastest-moving, responding to political events on scales of weeks to months. This natural hierarchy suggests that the CMH state space has a well-defined slow manifold parameterised primarily by demographic and long-run economic variables, with political and psychological variables approximately tracking their quasi-equilibrium values as functions of this slow backbone. Projecting the dynamics onto this slow manifold filters the high-frequency fluctuations that drive rapid trajectory divergence and yields a reduced system that is substantially more predictable on decadal time scales. The practical implication is that regime-level predictions conditioned on slow-manifold variables — asking whether a society will remain structurally stressed or structurally stable over the next twenty years, given its current demographic trajectory — are far less subject to the Lyapunov Wall than event-level predictions about specific political outcomes.

Complementing the slow manifold approach, Critical Slowing Down (CSD) theory provides a set of mathematically universal warning signals that a system is approaching a qualitative transition — a bifurcation, a tipping point, a regime shift. The theoretical foundation rests on a fundamental result of bifurcation theory: as a system's parameters approach a critical value at which a stable equilibrium disappears or changes qualitative character, the dominant eigenvalue of the system's Jacobian matrix at that equilibrium approaches zero. Mathematically, this means the system loses its ability to return to equilibrium after perturbations — perturbations decay more slowly, and the system's memory of its recent history extends further into the past. This phenomenon of critical slowing down produces three observable signatures in the time series of state variables: an increase in the variance of fluctuations around the equilibrium, an increase in the lag-one autocorrelation of those fluctuations, and a reddening of the power spectrum as low-frequency components gain relative power. These signatures are universal in the sense that they appear regardless of the specific mechanism driving the bifurcation — whether the approaching transition is a revolution, an economic collapse, or an ecological regime shift, the mathematical fingerprint of critical slowing down is the same.

The empirical record of CSD applications is encouraging. Scheffer and colleagues (2009) documented CSD signatures before transitions in lake ecosystems and in paleoclimate records. Dakos and colleagues (2012) demonstrated that CSD indicators preceded eight out of eight documented ecological regime shifts in long-term monitoring data. In the social domain, Guttal and colleagues (2016) found CSD signatures in economic time series preceding the 2008 financial crisis. For CMH, the most immediate application is to the slow-manifold variables already collected in the Arab Spring dataset: computing the rolling variance and autocorrelation of the SSI and its component variables in the years preceding observed transitions provides a preliminary test of whether CSD signatures are detectable in the available data. This test is implementable with the current dataset and constitutes a valuable methodological validation step that does not require the expanded historical panel needed for the Koopman programme.

The combination of slow manifold projection and Critical Slowing Down monitoring thus offers a practically immediate extension of CMH's predictive capacity at the regime-probability level. Its limitations are well-understood: CSD indicators provide advance warning that a transition is approaching, but they do not specify the type or direction of the transition, and they can produce false positives in systems with strong noise that mimics slowing-down behaviour. The lead time of CSD signals is also variable and difficult to quantify, ranging from a few years to a few decades depending on the rate at which the system's parameters approach the critical value. Within these limitations, the approach is deployable with the existing CMH dataset without waiting for the expanded historical panel that the Koopman programme requires.

3.3 Class III — Spectral Decomposition: Reservoir Computing and Koopman Operator / DMD

Where Class I and Class II approaches work within the constraints of the Lyapunov Wall — either by better characterising uncertainty at the event level or by shifting attention to slower regime dynamics — Class III approaches ask a different question altogether: not where the system will be, or what regime it will inhabit, but what enduring mathematical structures govern its long-run behaviour. These are features not tied to any particular trajectory and therefore not subject to the trajectory-level form of the Wall.

Reservoir computing is a machine learning architecture designed specifically for time-series prediction in non-linear and chaotic systems. The architecture consists of three components: a high-dimensional dynamical system with randomly fixed connections — the reservoir — whose internal state is driven by the input time series; a readout layer that maps the reservoir state to the desired output; and a training procedure that adjusts only the readout weights, leaving the reservoir itself unchanged. The reservoir acts as a computational substrate that implicitly projects the input signal into a high-dimensional space where non-linear relationships in the original signal become approximately linear and therefore separable by a simple linear classifier or regressor. The training procedure reduces to linear regression, which is computationally trivial and free of the vanishing-gradient problems that affect deep recurrent neural networks trained by backpropagation.

The theoretical foundations of reservoir computing relate it directly to the concepts of slow manifolds and Koopman operators discussed elsewhere in this section. The reservoir, when driven by a chaotic input signal, synchronises its internal dynamics to the low-dimensional attractor of that signal — a phenomenon known as generalised synchronisation. The reservoir state therefore provides a high-dimensional encoding of the input's recent history that is particularly rich in information about the slow, persistent components of the signal: precisely the components that survive the Lyapunov Wall. Pathak and colleagues demonstrated in 2018 that a modest reservoir computing system — a network of approximately one thousand neurons — could predict the evolution of the Lorenz system, the canonical example of deterministic chaos, for approximately eight Lyapunov times, roughly four times longer than classical numerical integration approaches. This result, replicated and extended in multiple subsequent studies, establishes reservoir computing as the current state of the art for data-driven prediction of chaotic systems in the regime where training data are abundant.

For Computational Macrohistory, the barrier to immediate application is data volume. Effective reservoir training requires several hundred to several thousand time steps of the target signal, with the exact requirement depending on the complexity of the dynamics and the size of the reservoir. The Arab Spring dataset provides thirteen annual observations per country — roughly two orders of magnitude below the minimum requirement. Even the expanded Western European historical panel described in Section 5, with approximately one hundred and forty usable years across ten countries, yields only on the order of one thousand country-year observations when pooled across the panel, which approaches but does not comfortably exceed the lower end of the feasible range. Reservoir computing is therefore categorised here as a medium-term candidate for CMH, contingent on the expansion of the empirical programme to cover longer historical periods and larger country samples. As the empirical database grows toward the fifty or more country-cases identified as the threshold for methodological expansion in the SSI-15 programme, reservoir computing becomes increasingly viable and should be revisited as a complement to the Koopman approach.

The Koopman Operator framework, implemented through Dynamic Mode Decomposition, is the most theoretically distinctive member of Class III and the leading candidate for integration with CMH. Its mathematical foundations and its connection to the axiomatic structure of CMH are the subject of the following section.

3.4 Summary Comparison

Table 1 summarises the comparative assessment. The Class column reflects the methodological classification developed above; the Temporal Horizon column reflects the prediction type at which each approach is assessed (event, regime, or spectral).

Table 1. Comparative assessment of candidate approaches, organised by methodological class.

Method	**Class**	**Pred. Level**	**Temporal Horizon**	**Applicability to CMH**	**Maturity**
Ensemble Methods	I — Uncertainty prop.	1 — Event	2–3× baseline	High	Very mature
Analog Forecasting	I — Uncertainty prop.	1 — Event	Limited by archive	Low (current)	Mature
Slow Manifold / CSD	II — Regime detection	2 — Regime	20–50 years	Very high	Mature
Reservoir Computing	III — Spectral decomp.	2–3	3–8× baseline	Medium	Recent
Koopman / DMD	III — Spectral decomp.	**3 — Spectral**	**50–100 yrs (hypothesis)**	**Very high**	**Developing**

Note on temporal horizon estimates: The figures in the Temporal Horizon column reflect the prediction type at which each approach is assessed, not uniform empirical benchmarks. For Class I methods, the "2–3× baseline" figure derives from ensemble spread analyses in atmospheric and economic forecasting applied to a CMH baseline of 5–15 years. For Class II methods, the "decades" estimate follows from the characteristic time scales of slow demographic and economic variables as formalised in CMH Axiom A3. For the

Koopman/DMD entry, the "50–100 yrs (hypothesis)" figure is a theoretical upper bound derived from the phase-accumulation formula in Section 2, assuming frequency estimation error of ±1% and structural stability over the forecast window; it is not an empirical result and is explicitly flagged as a working hypothesis pending the analysis of the companion paper.

4. THE KOOPMAN OPERATOR: THE LEADING CANDIDATE

Of the five approaches surveyed, the Koopman Operator framework stands out as the most promising for integration with CMH at the level of spectral structure persistence. Before describing how the framework works and how it connects to the CMH axiomatic system, it is important to be precise about the nature of the predictive content that Koopman modes may provide — and its limits — so as to avoid the confusion identified in Section 2 between description and prediction. The horizon of fifty to one hundred years cited in Table 1 for Koopman cyclic modes is a working hypothesis, not an empirical result: it is the theoretical upper bound on phase predictability given plausible frequency estimation errors, conditional on structural stability. Whether historical socio-political systems actually contain Koopman modes with the requisite stability properties is the empirical question that the programme described in Section 5 is designed to answer.

A Koopman mode with zero growth rate is a persistent oscillation: it neither grows nor decays, and its frequency can be estimated from data with increasing precision as the observation record lengthens. The prediction it supports is not that the system will be in a specific state at a specific future time — that remains subject to the full force of the Lyapunov Wall — but that the system will continue to oscillate with a certain period and amplitude. This predictive content is not subject to exponential decay, but to linear phase uncertainty and structural stability constraints: the phase estimate accumulates error linearly with the forecast horizon, and the prediction holds only so long as the structural parameters of the system remain within a stable range. This is a prediction at the spectral level, not the event level. It is falsifiable, because the predicted oscillation can fail to appear in out-of-sample data. It carries no implication about specific system states at specific future times.

4.1 The Core Idea: Making Non-Linear Systems Linear

The central mathematical challenge of forecasting in any non-linear system is that small differences in initial conditions lead to large differences in outcomes — the Lyapunov Wall. One natural response is to look for a transformation of the system that eliminates the non-linearity entirely. This is the strategy Bernard Koopman proposed in 1931, in a paper that remained largely theoretical for several decades before becoming computationally practical in the early twenty-first century.

The idea can be illustrated with an analogy. Consider the shadow of a rotating three-dimensional object projected onto a two-dimensional wall. The shadow's movement may appear irregular and hard to predict from the two-dimensional perspective: it contracts and expands, accelerates and slows, in patterns that seem to have no simple law. But if one knows the object is a sphere rotating at constant angular velocity, the three-dimensional motion is perfectly periodic and entirely predictable — the apparent irregularity is a consequence of the projection, not of the underlying dynamics. The Koopman approach does something mathematically analogous: it changes the representation space in which the system is described, moving from a low-dimensional non-linear description to a high-dimensional linear description. The irregular, apparently unpredictable motion in the original space becomes the projection of an exactly linear evolution in the extended space.

Technically, the Koopman operator acts not on the state of the system but on functions of the state, called observable functions. Let the system evolve according to the map Φ — so that if the system is in state x at time t, it will be in state $\Phi(x)$ at time $t + \Delta t$. The Koopman operator K associated with Φ is defined by the relation $K f = f \circ \Phi$: it maps any observable function f to the observable function whose value at state x equals the value of f at the future state $\Phi(x)$. This operation is linear in f regardless of whether Φ is linear or not: $K(\alpha f + \beta g) = \alpha Kf + \beta Kg$ for any functions f and g and scalars α and β. The remarkable consequence is that the non-linear dynamics of Φ, which are in general intractable, induce an exactly linear dynamics on the infinite-dimensional space of observable functions. Every non-linear system has a Koopman representation; the challenge is identifying a finite-dimensional approximation of this representation that is both tractable and informative.

The payoff of this linear representation is access to the full machinery of spectral analysis. A linear operator on a function space has eigenvalues and eigenfunctions that completely characterise its long-run behaviour. The eigenfunctions of the Koopman operator — functions φ_j such that $K \varphi_j = \lambda_j \varphi_j$ for a scalar eigenvalue λ_j — have particularly simple temporal evolution: if the system starts in state x_0, then the value of φ_j at time t is simply $\varphi_j(x_0)$ multiplied by λ_j^t. Any observable of the system can be expressed as a linear combination of these eigenfunctions, and its temporal evolution is therefore an exact superposition of exponential (or oscillatory) terms. For stable modes with $|\lambda_j| < 1$, the contribution decays; for neutrally stable modes with $|\lambda_j| = 1$, the contribution oscillates at constant amplitude — these are the persistent structures that support long-range spectral-level predictions. For unstable modes with $|\lambda_j| > 1$, the contribution grows, signalling an impending transition.

4.2 Modes, Frequencies, and the Three Prediction Levels

The spectral decomposition of the Koopman operator yields a set of dynamical modes, each characterised by a complex eigenvalue $\lambda_j = \rho_j$ e^{iω_j}, where ρ_j is the growth or decay rate and ω_j is the oscillation frequency. Any observable of the system — any measurable quantity that can be expressed as a function of the state — can be decomposed into contributions from these modes, each evolving independently in time. The connection to the three-level prediction hierarchy of Section 1 is not merely formal: it is structural.

Modes with growth rate $\rho_j < 1$ are transient: they decay exponentially and eventually become negligible relative to the persistent components. These modes dominate short-horizon event-level predictions. They capture the transient response of the system to perturbations — the rapid adjustments that follow a political shock, an economic crisis, or a demographic change — and their decay rate determines how quickly the system returns to its underlying attractor after such perturbations. Modes with growth rate $\rho_j > 1$ are unstable: their amplitudes grow over time, and their presence indicates that the system is moving away from a stable configuration. In the CMH context, modes with growing amplitude that are concentrated in the political and economic variables provide an early warning of impending structural transition — a signal that is closely related to, but formally distinct from, the Critical Slowing Down indicators of Class II. Modes with growth rate $\rho_j \approx 1$ are neutrally stable: their amplitudes neither grow nor decay, and they oscillate at frequency ω_j indefinitely, subject only to the constraint that the structural parameters θ of the system remain within a stable range. It is these neutrally stable modes that support spectral-level predictions beyond the event-prediction horizon.

Neutrally stable Koopman modes in historical systems correspond to cyclical patterns in the dynamics that are not driven by transient perturbations but are intrinsic to the system's structural organisation — patterns that would persist even in the absence of external shocks. Whether such patterns exist in macro-historical data is an empirical question; the qualitative historical literature provides theoretical motivation to expect them. Demographic cycles — the roughly twenty to thirty-year rhythm linking the relative size of young adult cohorts to episodes of political mobilisation and instability — have been proposed as structural features of agrarian and early industrial societies across multiple civilisations and time periods. The long-wave hypotheses of Kondratiev and Schumpeter, suggesting approximately fifty-year cycles of technological innovation and economic restructuring, have been debated in economic history for a century. Modelski's long-cycle theory of hegemonic succession posits approximately one-hundred-year cycles in the organisation of the international system. Whether these qualitatively proposed cycles correspond to genuine neutrally stable Koopman modes in quantitative historical data — rather than being post-hoc narratives imposed on noisy time series — is the empirical question that Section 5 is designed to address. The Koopman framework does not presuppose the answer; it provides a rigorous mathematical test of the hypothesis.

A critical qualification is necessary. The claim that neutrally stable Koopman modes support extended predictive horizons applies only under the assumption of structural stability: the parameter vector θ that characterises the system's structural environment must remain approximately constant over the forecast horizon. For a Koopman mode with period T estimated from data spanning an observation window W, the mode is predictive forward in time only so long as the system remains in the same structural regime. Major technological discontinuities — the introduction of mass communication, the digital revolution, a future transformation of comparable scope — fundamentally alter the structure of the dynamical system and render modes estimated from pre-discontinuity data inapplicable to post-discontinuity dynamics. This structural stability constraint is not unique to the Koopman approach; it is a general feature of all time-series modelling that assumes stationarity, and it is the empirical counterpart of the conditional nature of Axiom A2's ergodicity claim.

4.3 Dynamic Mode Decomposition: From Theory to Data

The practical implementation of Koopman operator theory is Dynamic Mode Decomposition, developed by Schmid in 2010 and placed on a rigorous theoretical foundation by Tu and colleagues in 2014. DMD takes as input a sequence of observed system states at successive time steps — exactly the kind of panel data that CMH collects — and extracts the dominant dynamical modes together with their frequencies and growth rates.

Mathematical notation preview

The following notation will be used throughout this paper and developed formally in the companion paper. Let $X(t) \in \mathbb{R}^n$ denote the CMH state vector at time t, where n is the number of observed variables. Standard DMD seeks the matrix A such that $X(t+1) \approx A\,X(t)$, then extracts the spectral decomposition of A — its eigenvalues and eigenvectors — to identify the dynamical modes. In the multi-task formulation, the same matrix A is shared across all countries c in the panel: $X^{(c)}(t+1) \approx A\,X^{(c)}(t)$, with country-specific amplitude coefficients estimated separately for each c. In the Hankel extension, the state vector is augmented with temporal lags to construct a Hankel matrix H whose columns are $[X(t), X(t+1), \ldots, X(t+\tau)]^T$ for a chosen delay τ, increasing the effective dimensionality of the observable space without requiring additional observations. These three components — standard DMD, multi-task estimation, and Hankel augmentation — combine into the multi-task Hankel-DMD estimator that is the computational core of the proposed empirical programme.

The algorithm finds the best linear map relating each observed state to the next, decomposes that map into spectral components, and identifies the modes capturing the most variance. The output is a set of spatial patterns — describing which variables participate in each oscillation and with what relative timing — paired with temporal signatures. Reconstruction of past behaviour and projection of future behaviour are then a matter of summing mode contributions weighted by initial amplitudes.

For CMH, a particularly relevant extension is the multi-task formulation: rather than fitting independent DMD models to each country separately — impractical with short time series — the modes and frequencies are estimated jointly across all countries simultaneously, on the assumption that countries in similar structural conditions share the same underlying dynamical architecture. This assumption is the empirical counterpart of CMH Axiom A2, as discussed below. A second extension, Hankel-DMD, addresses the limited temporal depth of annual panel data by constructing augmented data matrices that embed temporal lags, effectively increasing the informational content available to the decomposition without requiring additional observations.

4.4 Algorithmic Extensions for Historical Panel Data

The standard DMD algorithm, as developed by Schmid (2010) and Tu and colleagues (2014), is designed for densely sampled time series from a single trajectory. Its direct application to the CMH empirical setting — annual observations of five to twenty-five variables across a panel of ten to twenty countries, each country providing a separate but structurally related trajectory — requires three algorithmic extensions that together define the multi-task Hankel-DMD estimator proposed as the computational core of the Koopman-CMH programme.

The first extension addresses the limited temporal depth of the available time series. Standard DMD requires a number of time steps substantially larger than the dimension of the state space in order to estimate the propagator matrix A reliably. When the time series is short relative to the state dimension — as is the case for CMH data, where thirteen to two hundred annual observations must support the estimation of modes in a space of five to twenty-five variables — the standard algorithm is ill-conditioned and produces unreliable spectral estimates. The Hankel extension addresses this by constructing an augmented data matrix whose columns concatenate the current state with lagged values of the state at delays of one, two, and up to τ time steps. This procedure, motivated by the Takens embedding theorem (1981), which guarantees that a sufficiently rich delay embedding reconstructs the attractor of the generating system, effectively multiplies the informational content of each time step by a factor of $\tau + 1$. For CMH data with annual observations and modes of interest at periods of twenty to one hundred years, a delay of five to ten years is appropriate, yielding augmented state vectors of dimension five to two hundred and fifty depending on the number of variables and the delay.

The second extension addresses the multi-country structure of the CMH panel. Rather than estimating a separate propagator matrix for each country — which is infeasible given the short time series available per country — the multi-task formulation imposes a shared propagator across all countries in the panel, with country-specific initial amplitude coefficients. The shared propagator encodes the modes and their frequencies, which are the quantities of interest for the Koopman-CMH programme; the amplitude coefficients encode how strongly each mode is excited in each country's particular historical trajectory. This decomposition is directly motivated by Conjecture 1 (Section 4.5): if countries in the same structural regime share the same Koopman eigenfunctions, then a shared propagator is the correct statistical model, and pooling information across the panel dramatically improves the precision of the spectral estimates.

The third extension concerns the choice of observable functions. Standard DMD operates in the original state space, which corresponds to choosing the identity as the observable function. This choice is appropriate when the Koopman eigenfunctions lie approximately in the span of the state variables themselves — a condition that holds for systems close to a linear regime but may fail for strongly non-linear systems. Kernel DMD (Williams et al., 2015) extends the approach to a richer class of observable functions by implicitly working in a high-dimensional feature space defined by a kernel function, without requiring the explicit specification or computation of those features. For CMH, where the non-linear interactions among variables

— particularly the multiplicative interactions between demographic, economic, and political variables — are central to the dynamics, kernel DMD may be necessary to capture the relevant Koopman structure. The trade-off is computational: kernel methods scale poorly with dataset size, and their application to the historical panel will require careful regularisation to avoid overfitting. The companion paper will evaluate both the linear and kernel versions of the multi-task Hankel-DMD estimator and provide guidance on the conditions under which each is preferred.

4.5 Connection to the CMH Axiomatic Framework

The Koopman / DMD framework connects to the CMH axiomatic system at multiple points. The most important of these connections concerns Axiom A2 (conditional ergodicity), which postulates that systems operating under stable structural conditions belong to the same statistical class and can be modelled by the same invariant distribution. In the language of Koopman operator theory, this translates into a conjecture about shared spectral structure:

> **Conjecture 1 (Koopman–A2 Correspondence):**
>
> *Let $X^{(c)}(t)$ denote the state trajectory of country c under structural parameter vector θ satisfying the conditions of CMH Axiom A2 (conditional ergodicity) and Axiom A4 (dynamic continuity), and assume weak mixing and finite-variance noise conditions on the stochastic forcing. Then the leading Koopman eigenfunctions of the systems $\{X^{(c)}\}$ are approximately shared across the panel: the eigenvalues λ_j and the spatial structure of the dominant modes Φ_j are common to all countries in the same ergodic class, up to country-specific amplitude coefficients $b_j^{(c)}$. Empirical failure of this conjecture — modes that are not shared — would constitute evidence against A2 and require revision of the ergodicity assumption.*

This conjecture is not proved here; its proof requires the full mathematical apparatus of Koopman operator theory applied to stochastic systems, which is the subject of the companion paper. What can be stated here is the following. The conjecture is consistent with the existing CMH axioms: A2 guarantees that countries in the same regime share a common ergodic distribution, and ergodic distributions have well-defined spectral decompositions under the Koopman operator. The conjecture adds the specific claim that these spectral decompositions are shared at the level of the leading modes. It is falsifiable empirically through the multi-task DMD estimation: if the modes extracted jointly from the panel are a significantly worse fit than modes extracted independently for each country, the conjecture is rejected.

Beyond A2, the framework connects to A5 (endogenous indeterminacy) in a way that requires careful statement. A5 asserts that the trajectory of the CMH state vector is not fully determined by initial conditions and structural parameters — there exists an irreducible stochastic component that makes precise trajectory-level prediction impossible in principle. A potential objection to the Koopman programme is that the existence of stable spectral modes appears to contradict this indeterminacy: if neutrally stable cycles persist indefinitely, does the system not become, in some sense, predictable? The objection dissolves once the level of analysis is held explicit. A5 operates at the level of the state trajectory X(t): it denies that the system will be at a specific point in state space at a specific future time. The Koopman spectral structure operates at the level of the observable functions and their long-run statistical properties: it asserts that certain linear combinations of observables will oscillate with a characteristic frequency, not that the system will occupy a determinate state. Indeterminacy at the trajectory level is entirely compatible with regularity at the structural level — the two claims refer to different mathematical objects. Axiom A7 (predictive decay) is refined rather than contradicted: the event-level component decays exponentially as A7 states, but the spectral-level component does not, and this distinction is now made formally explicit within the CMH architecture. Axiom A8 (computational falsifiability) is satisfied directly: DMD produces out-of-sample reconstruction predictions whose accuracy can be measured against held-out data using standard probabilistic scoring rules.

5. AVAILABLE HISTORICAL DATASETS AND EMPIRICAL STRATEGY

A meaningful empirical test of the Koopman / DMD framework in a CMH context requires historical panel data meeting four criteria that follow directly from the mathematical requirements of the method. These criteria distinguish the appropriate dataset from the Arab Spring panel used in WP-2026-003 and WP-2026-004 and make explicit the conditions under which the proposed test is technically valid.

The first criterion is temporal depth. To identify a mode with period T using DMD, a minimum of three to five complete cycles must be present in the data, requiring a series length of at least 3T. For the demographic cycles hypothesised above (T ≈ 25 years), this implies a minimum series length of seventy-five years; for hegemonic cycles (T ≈ 100–150 years), a minimum of three hundred to five hundred years. In practice, two hundred years of annual observations — available for Western European countries from several sources — is sufficient to identify modes with periods up to approximately sixty years with adequate statistical confidence.

The second criterion is sampling frequency. Annual observations are appropriate for the demographic, economic, and political variables of the CMH state space; sub-annual frequency is unnecessary for dynamics operating on decadal time scales and introduces additional measurement noise.

The third criterion is dimensional coverage. The state vector must include variables from at least three of the five CMH state-space dimensions (D, E, P, S, Ψ) to allow meaningful multi-variable mode identification. The Ψ dimension (collective psychology) cannot be populated from available historical sources before approximately 1990; it is excluded from the initial test.

The fourth criterion is approximate stationarity within sub-periods. DMD requires that the dominant dynamics be approximately stationary over the estimation window — a condition that holds within structurally stable historical regimes but may be violated across major discontinuities such as the two world wars or the post-1945 institutional reconstruction. The estimation window should therefore be chosen to avoid straddling such discontinuities, or the analysis should be conducted separately on sub-periods.

Several existing datasets satisfy these criteria to varying degrees. Table 2 provides an overview of the most relevant sources, with an explicit stationarity assessment for each.

Table 2. Historical datasets suitable for Koopman / DMD analysis, with dataset selection criteria.

Dataset	**Coverage**	**Min. Series**	**Key Variables**	**CMH Dim.**	**Stationarity**
Polity V / Polity 5	1800–2018, ~170 countries	218 yrs	Regime score, DEMOC, AUTOC, DURABLE	P	Locally stationary by regime
V-Dem	1789–2023, ~180 countries	234 yrs	Liberal Democracy Index + 4 indices	P	Locally stationary by regime
Maddison Project	1820–2020 (some from 1300)	200+ yrs	GDP per capita, population	E	Non-stationary; log-differenced

Correlates of War	1816–present	200+ yrs	Conflict onset, intensity, type	P/S	Stationary (event counts)
Jordà–Schularick–Taylor	1870–2020, 17 countries	150 yrs	GDP, credit/GDP, inflation, asset prices	E	Mixed; requires transformation
Human Mortality DB / UN	1750–present (Europe)	270 yrs	Age structure, fertility, mortality	D	Slow trend; quasi-stationary

For a first empirical test, the most tractable combination is the V-Dem dataset (political variables, 1789 to 2023), the Maddison Project Database (economic output, 1820 to 2020), and the Correlates of War project (conflict variables, 1816 to the present), applied to the countries of Western Europe — France, the United Kingdom, Germany, Italy, Spain, the Netherlands, Sweden, and two or three others — over the period 1820 to 1913 and 1945 to 2020, excluding the two world war periods to preserve approximate stationarity. A note on the 1820–1913 sub-period is warranted: this window encompasses major structural discontinuities including the unification of Italy and Germany, the Franco-Prussian War, and the Paris Commune. These events may violate the stationarity assumption for modes with periods shorter than the inter-event interval. The practical implication is that the 1820–1913 sub-panel should be treated with caution for mode identification at periods below approximately thirty years, and that robustness checks comparing mode estimates across sub-windows within this period are essential before drawing conclusions. The 1945–2020 sub-panel, covering a longer period of relative institutional stability in Western Europe, is likely to be more informative for the identification of medium-frequency modes. Together the two sub-panels yield approximately one hundred and forty years of usable annual observations, eight to ten countries, and coverage of the P, E, and conflict-related S dimensions of the CMH state space. The D dimension can be added through historical reconstructions from the Human Mortality Database.

This dataset will be used to address the three foundational questions formalised in Conjecture 1: whether persistent Koopman modes exist in the empirical data, whether they are shared across the country panel, and whether their frequencies correspond to historically documented cycles. A positive answer to all three questions would constitute the first empirical support for the Koopman-CMH integration and would justify the full mathematical development of the companion paper. A negative answer would require revision of either the conjecture or its underlying axiomatic assumptions — which is equally a scientific contribution.

A practical note on missing data is warranted. Historical panel datasets of this kind are inevitably incomplete, particularly in the earlier decades of the observation window. V-Dem provides near-complete coverage for Western Europe from 1900 onward and reasonable coverage from 1820, with gaps concentrated in periods of state dissolution or occupation. The Maddison Project fills gaps through interpolation using adjacent-year observations and cross-country regression, with the interpolated values clearly flagged in the dataset documentation. For the DMD estimation, missing observations require either listwise deletion of affected country-years, multiple imputation, or matrix completion methods; the choice among these strategies affects the statistical properties of the mode estimates and will be evaluated systematically in the companion paper. As a conservative baseline, the empirical test will be conducted on the subset of country-years for which all selected variables are observed without interpolation, with sensitivity analyses extending to the imputed sample.

6. EXPECTED CONTRIBUTIONS AND LIMITATIONS

If the empirical programme described above succeeds in identifying persistent Koopman modes in the Western European historical panel, the Koopman-CMH integration will make several distinct contributions to the scientific programme of Computational Macrohistory.

At the level of practical application, a successful programme would extend the CMH predictive horizon from the current five to fifteen years to a range of fifty to one hundred years — not for event-level forecasts, which remain subject to the Lyapunov Wall, but for phase-level statements: estimates of where in a multi-decadal or century-scale cycle a given society is currently located, and what that position implies for the probability of structural transitions over the coming decades. This is a contribution at the second and third levels of the prediction hierarchy introduced in Section 1, not the first.

The theoretical contribution is no less consequential. Conjecture 1 formalises the connection between the Koopman spectral structure and CMH Axiom A2, transforming the ergodicity assumption from an untested theoretical postulate into an empirically falsifiable claim. If the modes are not shared across countries in similar structural conditions, A2 requires revision. This makes the empirical test simultaneously a test of the method and a test of the axiomatic framework — the kind of disciplinary self-scrutiny that Axiom A8 demands.

Worth noting separately, and applicable even if the full Koopman programme succeeds only partially, is the three-level prediction hierarchy itself as an organising framework for CMH forecasting. The distinction between event prediction, regime probability, and spectral structure persistence clarifies what CMH can and cannot claim at each temporal horizon, and provides a language for evaluating future methodological extensions without the ambiguity that has sometimes surrounded claims of "extended prediction" in the complex systems literature.

The limitations of this programme must be stated with equal clarity. The Koopman / DMD framework extends the predictive horizon for the cyclic component of historical dynamics only so long as the underlying dynamical structure — the frequencies and growth rates of the modes — remains approximately constant. Fundamental structural transformations, such as a technological revolution that introduces entirely new variables into the dynamics, will render modes estimated from historical data inapplicable. This is the same constraint that governs all CMH modelling through Axiom A2, and it is not specific to the Koopman framework.

Additionally, the identification of modes from historical data is subject to statistical uncertainty that grows linearly with the forecast horizon, even for neutrally stable modes. A frequency estimated to within one percent accuracy will accumulate a half-cycle phase error over approximately fifty years. Long-range projections are therefore best understood as distributional statements about the existence and approximate period of structural cycles, carrying credible intervals that widen with the horizon, rather than as deterministic predictions of specific future states.

Finally, the claim that Koopman modes provide prediction rather than post-hoc description rests entirely on out-of-sample validation. A mode identified from data up to time T is predictive only if it continues to appear in data after T and generates accurate phase forecasts for that later period. Operationally, this means: a mode with estimated period $\hat{T}$ and phase $\hat{\varphi}$ at time T_0 is considered predictive if, in held-out data covering [T_0 + k, T_0 + k + $\hat{T}$] for k > 0, the same spectral peak appears within a tolerance of ±10% of $\hat{T}$ and the phase error does not exceed half a cycle. This criterion is strict enough to distinguish genuine predictive content from retrospective pattern-matching, and it is the operationalisation of Axiom A8 in the Koopman context. It will be the primary standard of evidence in the companion paper.

7. CONCLUSION

This paper has examined whether and how the predictive horizon of Computational Macrohistory can be extended beyond the Lyapunov Wall that currently limits its quantitative forecasts to a range of five to fifteen years. The analysis confirms that the Wall cannot be eliminated — it is a structural feature of chaotic dynamics, not a technological obstacle — but that it can be partially circumvented by ascending a hierarchy of prediction types: from event prediction, where the Wall applies in full force, to regime probability, where the effective decay is substantially slower, to spectral structure persistence, where genuinely persistent structures may carry predictive content on horizons of decades or centuries.

Five candidate approaches were surveyed and organised into three methodological classes. Ensemble Methods and Analog Forecasting (Class I) improve the characterisation of uncertainty at the event-prediction level without pushing the Wall back. Slow Manifold theory and Critical Slowing Down (Class II) extend the horizon at the regime-probability level and are immediately applicable to the existing CMH dataset. Reservoir Computing and the Koopman Operator / DMD framework (Class III) target the spectral level and carry the greatest potential for genuinely long-range predictive content, at the cost of more demanding data requirements.

The Koopman Operator framework emerged as the leading candidate at the spectral level, owing to its natural connection with CMH Axiom A2 (conditional ergodicity) formalised in Conjecture 1, its compatibility with the panel structure of available historical datasets, and its capacity to produce falsifiable predictions about persistent cyclic structures. The empirical programme proposed here — applying multi-task Hankel-DMD to a Western European historical panel spanning approximately two centuries — is designed to test Conjecture 1 and to provide the first empirical basis for the Koopman-CMH integration.

The present paper is a methodological roadmap. The companion paper, currently in preparation, will develop the full mathematical treatment: the formal proof of the conditions under which Conjecture 1 holds within the CMH axiomatic system, the derivation of the Hankel-DMD estimator and its statistical properties in the panel setting, the results of the empirical test on the Western European data, and the implications for the CMH predictive architecture. Readers with a background in operator theory, stochastic dynamical systems, or spectral analysis of time series are directed to that paper for the technical foundations.

The broader ambition that motivates this programme is to establish Computational Macrohistory not only as a framework for medium-term probabilistic forecasting of socio-political instability — the primary contribution of WP-2026-001 through WP-2026-004 — but as an anticipatory science capable of characterising the slow structural rhythms of historical change on scales of decades and centuries. The prediction hierarchy introduced in this paper provides the conceptual architecture for that ambition, and the Koopman framework provides the most mathematically credible path toward realising it. The core insight is perhaps best stated directly: CMH does not extend prediction by overcoming chaos, but by redefining the objects of prediction under chaos.

REFERENCES

Crutchfield, J. P., & McNamara, B. S. (1987). Equations of motion from a data series. Complex Systems, 1(3), 417–452.

Dakos, V., Carpenter, S. R., Brock, W. A., Ellison, A. M., Guttal, V., Ives, A. R., ... & Scheffer, M. (2012). Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. PLOS ONE, 7(7), e41010.

Fenichel, N. (1979). Geometric singular perturbation theory for ordinary differential equations. Journal of Differential Equations, 31(1), 53–98.

Fontaise, G. (2026a). Computational Macrohistory: Axiomatic Foundations for a Quantitative Science of Large-Scale Social Systems. FICSS Working Paper WP-2026-001. Zenodo. https://doi.org/10.5281/zenodo.18288165

Fontaise, G. (2026b). Computational Macrohistory: Operational Framework — State Space, Variables, and Dynamic Equations. FICSS Working Paper WP-2026-002. Zenodo. https://doi.org/10.5281/zenodo.18646832

Fontaise, G. (2026c). Computational Macrohistory: Exploratory Empirical Application — The Arab Spring as a Preliminary Test Case for Structural-Demographic Theory. FICSS Working Paper WP-2026-003. Zenodo. https://doi.org/10.5281/zenodo.18848734

Fontaise, G. (2026d). Structural Stress and Political Instability in the MENA Region: A Computational Macrohistory Analysis of the Arab Spring. FICSS Working Paper WP-2026-004. Zenodo. https://doi.org/10.5281/zenodo.19661257

Gao, J., Cao, Y., Tung, W., & Hu, J. (2007). Multiscale Analysis of Complex Time Series: Integration of Chaos and Random Fractal Theory, and Beyond. Wiley.

Guttal, V., Raghavendra, S., Goel, N., & Hoarfrost, Q. (2016). Lack of critical slowing down suggests that financial meltdowns are not critical transitions, yet rising variability could signal systemic risk. PLOS ONE, 11(1), e0144198.

Goldstone, J. A., Bates, R. H., Epstein, D. L., Gurr, T. R., Lustik, M. B., Marshall, M. G., Ulfelder, J., & Woodward, M. (2010). A global model for forecasting political instability. American Journal of Political Science, 54(1), 190–208.

Jordà, O., Schularick, M., & Taylor, A. M. (2017). Macrofinancial history and the new business cycle facts. NBER Macroeconomics Annual, 31(1), 213–263.

Koopman, B. O. (1931). Hamiltonian systems and transformation in Hilbert space. Proceedings of the National Academy of Sciences, 17(5), 315–318.

Lorenz, E. N. (1963). Deterministic non-periodic flow. Journal of the Atmospheric Sciences, 20(2), 130–141.

Lorenz, E. N. (1969). Atmospheric predictability as revealed by naturally occurring analogues. Journal of the Atmospheric Sciences, 26(4), 636–646.

Maddison Project Database (2020). University of Groningen. https://www.rug.nl/ggdc/historicaldevelopment/maddison/

Marshall, M. G., Gurr, T. R., & Jaggers, K. (2019). Polity IV Project: Political Regime Characteristics and Transitions, 1800–2018. Center for Systemic Peace.

Medio, A., & Gallo, G. (1995). Chaotic Dynamics: Theory and Applications to Economics. Cambridge University Press.

Mezić, I. (2013). Analysis of fluid flows via spectral properties of the Koopman operator. Annual Review of Fluid Mechanics, 45, 357–378.

Pathak, J., Hunt, B., Girvan, M., Lu, Z., & Ott, E. (2018). Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. Physical Review Letters, 120(2), 024102.

Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., & Sugihara, G. (2009). Early-warning signals for critical transitions. Nature, 461(7260), 53–59.

Schmid, P. J. (2010). Dynamic mode decomposition of numerical and experimental data. Journal of Fluid Mechanics, 656, 5–28.

Takens, F. (1981). Detecting strange attractors in turbulence. In D. Rand & L.-S. Young (Eds.), Dynamical Systems and Turbulence, Lecture Notes in Mathematics (Vol. 898, pp. 366–381). Springer.

Tetlock, P. E., & Gardner, D. (2015). Superforecasting: The Art and Science of Prediction. Crown Publishers.

Tu, J. H., Rowley, C. W., Luchtenburg, D. M., Brunton, S. L., & Kutz, J. N. (2014). On dynamic mode decomposition: Theory and applications. Journal of Computational Dynamics, 1(2), 391–421.

V-Dem Institute (2024). V-Dem Dataset v14. University of Gothenburg. https://www.v-dem.net

Williams, M. O., Kevrekidis, I. G., & Rowley, C. W. (2015). A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition. Journal of Nonlinear Science, 25(6), 1307–1346.