Beyond the Wall

How Computational Macrohistory Learned to Ask Better Questions

The Lyapunov Wall cannot be demolished.
But there is more than one way to get to the other side.

A deep exploration of what comes after the hardest limit in the science of history

PART 1: WHERE WE LEFT OFF

The Promise That Was Broken

In a previous guide in this series — *The Lyapunov Wall* — we made a confession that most prediction systems are not willing to make.

Computational Macrohistory, for all its mathematical machinery and axiomatic foundations, cannot predict the distant future. Not because the theory is wrong. Not because we haven't collected enough data. Not because we need a faster computer. But because the systems we study — societies, economies, political regimes — are governed by a property of mathematics called sensitive dependence on initial conditions. Tiny differences in starting conditions amplify exponentially over time. The amplification is unstoppable. And once it reaches the scale of the system itself, any prediction based on the starting conditions becomes meaningless.

This is the Lyapunov Wall. It applies to weather (the reason meteorologists can't tell you what it will be like in two months), to turbulent fluids, to financial markets, and to the socio-political dynamics that CMH studies. For social systems, our best estimate places the Wall at roughly five to fifteen years for quantitative, specific forecasts. After that, the mathematics unmistakably runs out.

The guide ended there, on an honest note that some readers found frustrating. If the Wall is truly a wall — if no amount of genius, computing power, or data can break through it — why keep building? What is the point of a science of history that admits it cannot predict history?

That question deserves a full answer. This guide is that answer.

The answer, it turns out, is more interesting than the question.

A Question Within a Question

Here is the thing about the Lyapunov Wall that took years to fully appreciate: it is a limit on a *specific kind* of prediction. Not prediction in general. Not knowledge in general. A specific kind.

When we say "predict the future," we sound like we are asking one question. We are actually asking at least three different questions, stacked on top of each other, dressed up in the same words. And the Wall treats each of them very differently.

Question one: *Where exactly will this system be at time T?* This is what a navigator asks when plotting a course. This is what Hari Seldon was supposedly doing in Asimov's Foundation novels — calculating the precise position of galactic civilisation across millennia. This is also the question the Lyapunov Wall hits hardest. For social systems, it hits it within fifteen years. Anything beyond that is mathematical noise.

Question two: *What kind of place will this system be in, roughly, over the next few decades?* Not the precise coordinates — the type of terrain. Stable or unstable? Approaching a crisis or pulling back from one? Accumulating structural stress or releasing it? This is a coarser question, and a fundamentally different one. The Wall applies here too, but its force is substantially weaker. Horizons of twenty, thirty, even fifty years become plausible.

Question three: *Does this system have persistent rhythms — slow, deep oscillations that recur across generations — and what are they?* Not where the system will be, and not even what kind of place it will be in, but whether there are patterns that repeat regardless of the specific trajectory. Seasons, not individual days. Tides, not individual waves. For this third question, the Lyapunov Wall has almost nothing to say.

These three questions define what we now call the CMH prediction hierarchy. They are the conceptual core of our new working paper, *Beyond the Lyapunov Wall* — the fifth in the FICSS series. And understanding them changes everything about how to think about what a science of history can and cannot deliver.

PART 2: THE PREDICTION HIERARCHY

Why the Same Question Isn't

Before we get into the mathematics, let me try to make the distinction between the three levels vivid with a series of analogies, because this is the single most important idea in this guide.

Think about water.

If I ask you to predict the exact path of a single water molecule flowing down a river — when it will be on the left bank, when it will be in a turbulent eddy, when it will surface and when it will be submerged — you would be right to laugh. Individual water molecules in a turbulent river behave chaotically. They are as unpredictable, over the medium term, as the billiard balls we discussed in the previous guide. This is the first-level question, and the Lyapunov Wall demolishes it.

If I ask you to predict whether a river will be flooded or dry in a given month, based on seasonal rainfall patterns and upstream reservoir levels, you have a fighting chance. You can make reasonable probabilistic estimates. You can say "flood probability in February is 60%, drought probability is 15%." This is a regime-level question — the second level — and it is substantially more tractable.

But now suppose I ask you: does this river flood roughly every seven years? Has it done so consistently for the past two centuries? The answer to that question — which is a question about the *structure of the dynamics*, not about any particular trajectory or regime — requires only historical observation and statistical analysis. It is not threatened by the chaos of individual molecules at all. The seven-year cycle doesn't care what any individual water molecule does. This is the third-level question, and it has a radically different relationship to chaos.

Or think about music. If you hear a symphony orchestra play a piece you have never heard before, you cannot predict the next note with certainty. That's the first-level question. You might be able to say "this sounds like it's building toward a dramatic climax, rather than winding down quietly" — that's the second-level, regime question. But you can absolutely tell, within a few minutes, whether the piece is in a major or minor key, whether it is following sonata form, and roughly what harmonic patterns structure it. That's the third-level question — the spectral structure of the music — and it is entirely accessible even though the specific sequence of notes was unpredictable.

History is a symphony we are hearing for the first time. We cannot predict the next note. We might be able to sense whether it is building to a climax or resolving to stillness. And we may — this is the wager — be able to identify its key signature.

Level One: Event Prediction

The first level of the hierarchy is what most people mean when they talk about predicting history. "When will the next revolution happen?" "Will Country X remain stable over the next five years?" "What is the probability of a major financial crisis within the next decade?"

These are event-prediction questions. They are trajectory questions. They ask where the system will be at a specific future point, with enough specificity to distinguish different outcomes from each other. They are what the previous four papers in the CMH series have been building toward — and they are what the Lyapunov Wall governs directly.

At this level, CMH's Systemic Stress Index (SSI), tested across eleven countries in our Arab Spring study, demonstrates that useful probabilistic statements can be made over horizons of roughly five to ten years. The SSI correctly distinguished countries that experienced revolutionary transition from those that remained stable in 71.4% of pairwise comparisons. This is meaningful information. But it is meaningful only within the temporal range where the Wall hasn't yet demolished the signal.

Beyond ten to fifteen years, event-level prediction degrades rapidly and eventually becomes meaningless. This is not a failure of the framework. It is the framework being honest.

Level Two: Regime Prediction

The second level asks a coarser but still informative question: what dynamical *regime* will the system be in, over a future interval? Not exactly where it will be — in which attractor basin, approaching which tipping point, in what broad structural configuration.

Think of the difference between asking "will it rain at 3pm on Tuesday?" (level one — trajectory) and "will this winter be wetter or drier than average?" (level two — regime). The second question is answerable with far more confidence and over far longer horizons, precisely because it doesn't ask about specific moments or specific trajectories. It asks about the statistical character of the dynamics.

For social systems, regime prediction focuses on the slow-moving structural variables — demographic pressure, inequality trajectories, institutional capacity — that change on timescales of decades rather than years. These variables have their own, much lower Lyapunov exponents. They are the deep currents of the system, not its surface waves.

CMH's demographic variables are a clear example. If we observe today that 32% of a country's population is between the ages of 15 and 24 — what structural-demographic theorists call a "youth bulge" — we can say with near-certainty that in fifteen years, those people will be between 30 and 39. They don't disappear. They don't age backward. The demographic trajectory of this cohort is almost entirely predictable, because demography is a slow variable with a very low Lyapunov exponent. What those people will *do* — whether they will find work, emigrate, become politically active, or organise in opposition to the state — depends on fast variables we cannot predict. But the structural pressure they will represent is already legible in current data.

This is regime prediction: not "what will happen" but "what structural conditions will obtain." And it is achievable over horizons the Wall barely touches.

Level Three: Spectral Structure

The third level is the most surprising, and the most powerful. It asks: does this system have persistent oscillations? Does it repeat? Not in a mechanical, clockwork way — social systems are not clocks — but in a statistical sense? Is there a pattern that tends to recur, even if the specific episodes within the pattern vary enormously?

The analogy I keep returning to is the tide.

You cannot predict, by examining the water molecules in any particular cubic metre of ocean, when the tide will turn. Individual water molecules are chaotic; their trajectories are unpredictable. But tides are predictable centuries in advance — not because we can see individual molecules, but because we understand the gravitational architecture that drives them. The moon's orbit is not chaotic. The way ocean basins resonate with tidal forcing is not chaotic. The cycle exists at a level of description entirely above the chaos of individual trajectories.

Do societies have tides? That is the empirical question.

The qualitative historical literature has long suggested that they do. Demographic cycles — periods in which a large young adult cohort enters a labour market and political system that cannot absorb them, followed by instability, followed eventually by resolution — have been identified across contexts ranging from ancient Rome to early modern Europe to the contemporary Middle East. Long waves of economic expansion and contraction — sometimes called Kondratiev waves — have been proposed with periods of fifty to sixty years. Cycles of political legitimacy — in which institutions accumulate authority, then exhaust it, then reconstruct it — have been discussed on timescales of forty to eighty years.

These cycles are not accepted by all historians. They are contested, imprecisely defined, and often identified with the benefit of hindsight. But they are suggestive of something real: the possibility that beneath the chaotic surface of historical events, there are structural oscillations that recur with enough regularity to be scientifically characterised.

The question is how to find them, if they exist. And this is where a piece of nineteenth-century mathematics, made computationally practical only in the last decade, enters the story.

PART 3: THE MAN WHO MADE CHAOS LINEAR

Bernard Koopman's 1931 Idea

On a Tuesday morning in 1931, a young mathematician named Bernard Osgood Koopman published a short paper in the *Proceedings of the National Academy of Sciences*. The paper was four pages long. It was titled "Hamiltonian systems and transformation in Hilbert space." Almost nobody read it.

This is unfortunate, because Koopman had just proposed one of the most conceptually elegant ideas in the history of dynamical systems: a way to make non-linear systems linear.

Let us spend a moment understanding why that matters so much.

The entire problem with chaotic systems — the reason the Lyapunov Wall exists — is that they are non-linear. In a linear system, small differences stay small. Errors don't amplify. Trajectories that start close together stay close together. Linear systems are mathematically tractable, predictable, and well-understood. In a non-linear system, small differences can become large ones. Trajectories diverge. The Wall appears.

So the dream would be: what if you could find a way to *represent* a non-linear system as if it were linear? Not by approximating it — not by linearising around a fixed point, as physicists routinely do — but exactly. What if there were some transformation of the system that preserved all the dynamics while eliminating the non-linearity?

Koopman proved that such a transformation exists. Always. For any dynamical system, no matter how complex or chaotic.

The trick is to stop thinking about the *state* of the system — the specific values of all its variables — and start thinking about *functions of the state*. Call these functions "observables." Instead of tracking where the system is, track quantities you can compute from where the system is: the temperature, the pressure, the population ratio, the product of two variables, the logarithm of a third.

The remarkable mathematical fact — the one Koopman proved in those four pages — is that for any non-linear system, no matter how wild its trajectories, the evolution of observables in time is *exactly linear*. The non-linearity has not disappeared. It has been relocated. It is still there, embedded in the geometry of the space of observables. But the dynamics — the way observables evolve through time — are linear.

What Linearity Gives You

Linear systems have a complete spectral theory. You can break them apart into independent components — modes — each characterised by a frequency and a growth or decay rate. These modes are the eigenfunctions of the linear operator governing the system. (If that sentence means nothing to you, don't worry — the idea is what matters, not the vocabulary.)

Think of it this way. A musical chord sounds complex when you hear it as a whole. But a signal-processing algorithm can decompose it into its constituent frequencies — the fundamental note, the harmonics, the overtones. Each frequency is independent. Each has its own amplitude. Together they produce the complex sound you hear, but separately they are simple, characterisable, predictable.

Koopman's framework does exactly this for dynamical systems. It decomposes the system's dynamics into independent modes, each with its own frequency and its own growth or decay rate.

Some modes have *positive* growth rates: they amplify over time. These are the modes associated with instability, with critical transitions, with systems that are about to blow up. Detecting them early is equivalent to detecting the early warning signals of approaching crises — the same signals discussed in our previous guides.

Some modes have *negative* growth rates: they decay over time. These are the transient features of the dynamics — the fingerprints of past events that gradually fade. They are relevant for event-level prediction over short horizons.

And some modes have *zero* growth rate: they neither amplify nor decay. They persist. These are the tides.

These neutrally stable Koopman modes — if they exist in a given system — are persistent oscillations. They recur with a characteristic period. They are not subject to the Lyapunov Wall in the same way as event-level predictions, because their predictive content is not about the system's location in state space, but about the structure of the space itself. The phase of an oscillation accumulates uncertainty linearly with time — not exponentially. A mode with an estimated period of twenty-five years, if it is real, will still be identifiable and approximately phase-predictable fifty years from now.

This is the spectral structure that the third level of the prediction hierarchy is trying to reach. And the Koopman framework is the mathematical tool that, in principle, allows us to get there.

PART 4: FROM THEORY TO DATA — THE DECOMPOSITION MACHINE

Dynamic Mode Decomposition

Bernard Koopman's 1931 paper was beautiful but impractical. To compute the eigenfunctions of the Koopman operator exactly, you would need to know the system's dynamics in closed form — which defeats the purpose, since the whole problem is that we don't.

The situation changed in 2010, when a fluid dynamicist named Peter Schmid published a paper introducing an algorithm he called Dynamic Mode Decomposition, or DMD. Schmid was not thinking about history. He was thinking about turbulent fluid flows — wind tunnels, jet exhausts, atmospheric dynamics. But what he created was a general-purpose tool for extracting Koopman modes from data, without needing to know the underlying equations.

Here is the core idea, stripped to its essentials.

Suppose you have a sequence of observations of a system at regular time steps. You measure it today, then tomorrow, then the day after, and so on. DMD looks at consecutive pairs of observations — the state of the system at time t and its state at time $t+1$ — and asks: what is the best linear map that transforms one into the other? It finds that map, decomposes it spectrally, and extracts the dominant modes — the patterns that explain the most variance in the data.

For CMH, the "observations" are the components of the state vector — the 25 variables spanning demographics, economics, politics, social structure, and collective psychology — measured annually across a panel of countries. The state at year t in country c is a list of 25 numbers. DMD processes those lists, year by year and country by country, and extracts the patterns that dominate the dynamics.

The output is a set of modes: spatial patterns (which variables move together, and in what proportions) paired with temporal signatures (at what frequency does this mode oscillate, and is it growing, decaying, or persisting?). If a neutrally stable mode exists in the data — a genuine, persistent oscillation that neither grows nor decays — DMD will find it. If no such mode exists, DMD will tell you that too.

The Multi-Task Trick

There is a practical problem. DMD works best with long time series — ideally, many hundreds of time steps per variable. For annual data on a single country, we are lucky to have sixty or seventy years of reliable measurements across all the required variables. Sixty time steps is not nearly enough to identify oscillations with periods of twenty, thirty, or fifty years.

The solution is to analyse multiple countries simultaneously — what statisticians call a multi-task approach. If France, Germany, the United Kingdom, Spain, Italy, and five other Western European countries share the same underlying dynamical modes — the same tides, as it were — then we can estimate those modes jointly, pooling across all ten countries at once. Ten countries times sixty years gives us six hundred observations, which is a much more reasonable basis for spectral estimation.

But this approach only makes sense if the assumption is actually true: if the countries really do share the same underlying modes. This assumption — that societies operating under similar structural conditions belong to the same statistical class, with the same characteristic frequencies — is not obvious. It might be false. France might oscillate on completely different timescales from Germany, for reasons specific to their respective histories, institutions, and geographies.

This is where the CMH axiomatic framework becomes directly relevant. Axiom A2 of the CMH axiomatic system — what we call the principle of conditional ergodicity — states precisely that systems operating under similar structural conditions tend to exhibit the same statistical regularities. It is the mathematical formalisation of the idea that history repeats itself across comparable contexts — not in its details, but in its structural dynamics.

The multi-task Koopman programme is, in effect, an empirical test of Axiom A2. If the modes are shared across the panel, A2 is supported. If they are not — if France and Germany turn out to oscillate on entirely different timescales with no shared structure — then A2 requires revision, and the programme will have discovered something important about the limits of the comparative method in historical analysis.

Either outcome is scientifically valuable.

The Conjecture

The most technically substantive contribution of *Beyond the Lyapunov Wall* is a formal conjecture — a precise, testable claim that the entire Koopman-CMH programme depends on. We call it the Koopman–A2 Correspondence.

In plain language, the conjecture says: *if CMH Axiom A2 holds — if countries in the same structural class share a common statistical distribution — then they also share the same leading Koopman modes. The frequencies and spatial structures of the dominant oscillations are common across the panel; only the amplitudes and phases differ from country to country.*

This is not proved here. Its proof requires the full mathematical apparatus of Koopman operator theory applied to stochastic systems — a treatment that will appear in the companion paper currently in preparation. What the conjecture does is make the shared-modes assumption explicit and falsifiable. It tells us exactly what we are betting on, and exactly what empirical evidence would prove us wrong.

A conjecture that cannot be falsified is not science. The Koopman–A2 Correspondence is falsifiable in a very concrete way: if the modes extracted jointly from the European panel are a significantly worse fit than modes extracted independently for each country, the conjecture is rejected. The companion paper will report the result, whatever it is.

PART 5: FIVE TOOLS FOR CLIMBING THE WALL

A Survey of the Landscape

The Koopman framework is the centrepiece of the new research programme — but it does not stand alone. *Beyond the Lyapunov Wall* surveys five distinct methodological approaches that can, in different ways and to different degrees, extend CMH's predictive range. They are not competitors; they are tools suited to different levels of the prediction hierarchy, and the best CMH programme will eventually use all of them. Four of them set the scene. The fifth — Koopman-DMD — is where the most theoretically ambitious work begins.

Tool One: Ensemble Methods

The first tool is the most widely deployed in operational forecasting, and the most familiar from meteorology: ensemble simulation.

Rather than running a single simulation from a single set of initial conditions, ensemble methods run dozens or hundreds of simulations, each starting from slightly different initial conditions drawn from the range of measurement uncertainty. The result is not a single predicted trajectory but a *distribution* of trajectories — a picture of the space of possible futures, weighted by probability.

If the ensemble is tightly clustered — all simulations producing similar outcomes — the system is in a predictable configuration. If the ensemble rapidly diverges into very different futures, the system is intrinsically unpredictable at this moment. The spread of the ensemble is itself informative.

Ensemble methods do not break the Lyapunov Wall. They characterise it more precisely. They do not extend the event-prediction horizon; they describe the uncertainty around that horizon more honestly. For CMH, they are an essential component of the simulation architecture for level-one prediction — the Prime Radiant computational pipeline that is the practical endpoint of the empirical programme.

Tool Two: Analog Forecasting

The second tool has a beautiful conceptual logic. Proposed by Edward Lorenz himself — the same Lorenz whose butterfly effect we discussed in the previous guide — analog forecasting asks: has the system been in a state similar to the current one before? If so, what happened next?

The idea is to search the historical record for "analogs" — past moments that closely resemble the present — and use the subsequent history of those past moments as a forecast for the present. No model required. No equations. Just pattern-matching in time.

For short time series and low-dimensional systems, analog forecasting can work remarkably well. It is essentially what experienced historians do when they say "this reminds me of the situation in 1930s Germany" or "this is similar to the conditions that preceded the French Revolution." The difference is that CMH does it quantitatively, systematically, and with explicit uncertainty quantification.

The limitation is dimensionality. For a 25-variable system like CMH's full state space, finding genuinely close analogs in the historical record requires an archive that grows exponentially with the number of variables. With the data currently available — a few decades of reliable annual data across perhaps a hundred countries — the archive is far too sparse to support reliable analog forecasting in the full 25-dimensional space.

Analog forecasting is a future tool, contingent on the expansion of the empirical programme to the N ≥ 30 country-cases required for statistical validation.

Tool Three: Slow Manifold Theory and Critical Slowing Down

The third tool operates at the second level of the prediction hierarchy — regime prediction — and it is the most immediately deployable.

Slow manifold theory, developed rigorously by mathematicians working on singular perturbation problems, formalises the separation between fast and slow variables that is already built into the CMH state space. The demographic variables of CMH — population structure, youth cohort size, urbanisation — evolve on timescales of fifteen to thirty years. Economic structural variables evolve on five to fifteen years. Political event variables evolve on months to a few years.

When you project the full system's dynamics onto the slow manifold — the subspace defined by the slow variables — you filter out the high-frequency, chaotic fluctuations that drive rapid trajectory divergence. What remains is a reduced, slower system that is intrinsically more predictable. The Lyapunov exponent of the slow manifold is smaller than that of the full system. The Wall is further away.

This is exactly how CMH's current empirical work on the Arab Spring proceeds: the Systemic Stress Index focuses on slow and medium variables (youth bulge, inequality, unemployment, regime stability, state capacity) rather than fast variables (public mood, media narratives, specific political events). By focusing on the slow manifold, the SSI achieves predictive relevance over the five-to-ten-year range that fast-variable models cannot reach.

Complementing the slow manifold approach, Critical Slowing Down theory provides a set of mathematically rigorous early warning signals that a system is approaching a qualitative transition. As a system nears a bifurcation — a tipping point where its qualitative character is about to change — its recovery time from small perturbations lengthens, its variance increases, and its internal correlations rise. These signatures have been empirically documented in ecosystems, climate records, and financial markets. They are universal: they apply regardless of the specific mechanism driving the transition.

In practice, this means that before a political crisis, the CMH variables should show increasing variance in public opinion indicators, longer recovery times from shocks (a scandal that used to blow over in two weeks takes two months; an economic dip that used to self-correct in a quarter takes two years), and rising correlations between indicators that were previously moving independently. When these signals fire, the system is approaching a tipping point — even if we cannot say which tipping point, or in what direction.

The combination of slow manifold projection and Critical Slowing Down monitoring is mature, theoretically justified, and immediately applicable to the existing CMH dataset. It is the single most practical extension of CMH's current capabilities in the near term.

Tool Four: Reservoir Computing

The fourth tool is the most exotic, and the furthest from deployment.

Reservoir computing is a form of machine learning specifically designed for predicting chaotic time series. It works by projecting the input signal into a high-dimensional space — the "reservoir," a large network of randomly connected computational units — where non-linear relationships become approximately linear. Only the output layer is trained; the reservoir itself is fixed and random. This simple architecture, paradoxically, produces powerful predictions.

In 2018, a team at the University of Maryland demonstrated that a reservoir computer could predict the Lorenz system — the canonical chaotic attractor — for approximately eight Lyapunov times, roughly four times longer than classical methods. The reason is that the reservoir implicitly learns a representation of the system's slow manifold — the part of the dynamics that is more predictable than the full trajectory.

For CMH, the barrier is data. Effective reservoir training requires several hundred time steps of reliable data. The Arab Spring dataset covers thirteen years annually across eleven countries —

far too sparse. Reservoir computing becomes available only when the empirical programme has expanded sufficiently. It is a medium-term candidate, not a near-term tool.

Tool Five: Koopman Operator and Dynamic Mode Decomposition

The fifth tool is the one we have already explored in depth, and the one *Beyond the Lyapunov Wall* identifies as the most theoretically promising for CMH's long-range ambitions.

Its advantages over the other approaches are worth stating explicitly.

Unlike ensemble methods and analog forecasting, Koopman-DMD operates at the third level of the prediction hierarchy — spectral structure — rather than the first or second. It is not trying to improve trajectory prediction. It is trying to identify the geometric architecture of the dynamics: the oscillations that persist regardless of which specific trajectory the system follows.

Unlike slow manifold theory, it does not require *a priori* knowledge of which variables are slow and which are fast. The DMD algorithm identifies the modal structure from the data itself, without theoretical presuppositions about timescale separation.

And unlike reservoir computing, it has a clear connection to the CMH axiomatic framework through the Koopman–A2 Correspondence conjecture, making its empirical test simultaneously a test of the method and a test of the theory. Few scientific programmes offer this kind of double leverage.

PART 6: THE DATA — TWO CENTURIES OF EUROPEAN HISTORY

Where to Begin Looking

The Koopman-CMH programme is not a purely theoretical exercise. It makes a concrete empirical prediction — that persistent, shared oscillatory modes exist in the dynamics of structurally comparable societies — and it intends to test that prediction against real historical data.

The test requires a specific kind of dataset, and the requirements follow directly from the mathematics. To identify a mode with period T reliably, you need at least three to five complete cycles in the data. For demographic oscillations hypothesised to span twenty-five years, that means at least seventy-five years of observations — and for the longer cycles some theorists have proposed, proportionally more. Two hundred years of annual data, available for most Western European countries from multiple archival sources, is sufficient to identify modes with periods up to roughly sixty years.

Annual sampling frequency is the right resolution for dynamics operating on decadal timescales: coarser measurements would miss important variation, while higher frequency would only add noise. The state vector must span at least three of CMH's five dimensions — demographics, economics, politics, social structure, and collective psychology — for multi-variable mode identification to be meaningful. The psychological dimension cannot currently be measured reliably for historical periods before roughly 1990 and will be excluded from the initial test. And the estimation window must avoid straddling major structural discontinuities — wars, revolutions, institutional collapses — where the system's fundamental character changes so abruptly that any spectral modes identified across the break would be artefacts of the discontinuity rather than features of the underlying dynamics.

The Western European Panel

With these criteria in mind, the most tractable initial dataset is a Western European historical panel, covering eight to ten countries — France, the United Kingdom, Germany, Italy, Spain, the Netherlands, Sweden, and two or three others — over two sub-periods: 1820 to 1913 and 1945 to 2020.

The two-world-war period is excluded for stationarity reasons. The European dynamics of 1914–1945 involve discontinuities so severe — state dissolution, mass mortality, territorial reorganisation, institutional collapse and reconstruction — that any spectral modes identified from that period are unlikely to represent the underlying peacetime dynamics.

The two sub-panels together yield approximately 140 years of usable annual observations, eight to ten countries, and coverage of the political (P), economic (E), and conflict-related social (S) dimensions. The political variables come from the V-Dem dataset, which provides continuous annual measurements of democracy indices and related institutional indicators back to 1789 for most European countries. The economic variables come from the Maddison Project Database, which provides GDP per capita and population estimates back to 1820. Conflict variables come from the Correlates of War project, which covers militarised interstate disputes and internal conflicts from 1816 to the present.

The demographic dimension (D) can be added through historical reconstructions from the Human Mortality Database, which covers age-structured mortality and population data for most European countries from the eighteenth century onward.

A Note on What We Expect to Find

It would be dishonest to pretend that we are entering this empirical programme without hypotheses. We are not.

The structural-demographic literature — associated most prominently with the work of Jack Goldstone, Randall Collins, and their collaborators — has long pointed to demographic cycles as a primary driver of political instability. The specific mechanism involves the size of the young adult cohort relative to the absorptive capacity of the economy and the political elite. When the youth cohort is large and opportunities are scarce, structural pressure builds. When it is small, or when expansion creates abundant opportunities, the pressure releases.

This mechanism generates a characteristic oscillation with a period loosely estimated at twenty to thirty years — the time it takes for a large young adult cohort to age through the period of peak political engagement and either find accommodation within the system or transform it through conflict. If this oscillation is real and recoverable by DMD, it will appear as a neutrally stable Koopman mode with a period in this range, characterised by coordinated movements in the D_2 (youth bulge) variable, the E_2 (inequality) variable, and the P_1 (regime stability) variable.

Whether this is actually what the data show — whether the oscillation is real, whether it is shared across the European panel, whether it is neutrally stable rather than slowly decaying — is precisely what the empirical programme will determine.

A negative result — no shared persistent modes, or modes that are not neutral but decaying — would be a genuine scientific finding. It would suggest that historical dynamics are chaotic "all the way down," with no recoverable spectral structure at the level of the available data. That would not make CMH wrong; it would make a particular extension of CMH wrong, and clarify the boundaries of the framework with greater precision than we currently have.

PART 7: WHAT CHANGES, AND WHAT DOESN'T

The Honest Ledger

Before the excitement of new mathematical tools and historical datasets carries us too far, it is worth sitting down and writing the honest ledger. What does the Koopman-CMH programme actually deliver, if it succeeds? What does it not deliver? And what remains exactly as it was before?

The Lyapunov Wall does not move. This bears repeating, because it is the most important thing to say. Event-level prediction — where exactly will Country X be in fifteen years — remains limited to the same five-to-fifteen-year range that the previous papers established. The Koopman framework changes nothing about this. If someone reads this guide and concludes that CMH now has a way to predict specific revolutions or regime changes decades in advance, they have misread it.

What changes is the range of legitimate scientific questions CMH can engage with. A successful Koopman-CMH integration, with persistent shared modes confirmed empirically, would allow a new kind of statement: not "Country X will experience political crisis within eight years" but "Country X currently appears to be in the late accumulation phase of a demographic-structural oscillation whose historical pattern, across comparable societies, has preceded periods of elevated systemic stress." That is not a trajectory prediction. It is a structural reading.

The difference between these two types of statement is not merely rhetorical. It is the difference between asking where a specific wave will crest and asking whether we are approaching high tide. The first question is largely unanswerable beyond the Lyapunov horizon. The second question might be answerable for decades.

What Prediction Means at the Spectral Level

Let us be precise about what "prediction" means for Koopman modes, because there is a real risk of confusion here.

A Koopman mode with period $\hat{T}$ and phase $\hat{\varphi}$ estimated at time T_0 generates a spectral-level prediction: that the same oscillation, with the same period (within a tolerance of ± 10%), will continue to appear in data after T_0, and that its phase will evolve consistently with the estimated frequency. This prediction is falsifiable: if the oscillation disappears in held-out data, or if its period changes substantially, the prediction is wrong.

Crucially, this prediction does not specify where the system will be at any particular future time. It specifies the *geometric structure of the space the system inhabits* — the frequency and pattern of the underlying oscillation. This is a prediction about the music, not about the next note.

The horizon over which phase predictions remain useful is bounded not by the Lyapunov exponent but by the precision of the frequency estimate. If the period of an oscillation is estimated as 25 years with a 2% uncertainty, the phase error accumulates at a rate of approximately one year of phase per 25 years of forecasting — linear, not exponential. After 50 years, the phase uncertainty is roughly 2 years; after 100 years, 4 years. This is manageable, and it represents a qualitatively different relationship to the future than anything event-level prediction can offer.

The important caveat — stated clearly in the paper — is that this calculation assumes the oscillation is genuinely neutral: neither growing nor decaying. If the mode has a small negative growth rate (slowly decaying), the prediction degrades faster than the linear calculation suggests. And if the structural parameters of the system change fundamentally — if, say, a technological revolution introduces entirely new variables into the dynamics — the modes estimated from historical data may no longer apply. This constraint is shared by all CMH modelling and is formalised in Axiom A2: it is the assumption that the structural class of the system is stable over the estimation window.

The ECG Analogy

The most useful analogy I have found for explaining what spectral-level prediction can and cannot do is the electrocardiogram.

An ECG does not tell you when you will have a heart attack, or whether you will have one at all. It does not predict specific events. What it does is characterise the electrical structure of your cardiac dynamics: the dominant rhythms, their regularity, the presence or absence of anomalous patterns. A cardiologist reading an ECG is operating at the spectral level: identifying whether the oscillatory structure of the heart is consistent with health or associated with elevated risk. "Your QRS complex shows this pattern" is not a trajectory prediction. It is a structural reading.

CMH, with a successful Koopman integration, aspires to be able to read a society's ECG. Not "this country will have a political crisis in 2031" but "this country's demographic-structural oscillation is currently in the phase that historically correlates with elevated instability risk over the coming decade." The two statements are different in kind, not just in degree. The first is a trajectory prediction, vulnerable to the Lyapunov Wall. The second is a spectral reading, subject to much slower degradation.

The cardiologist cannot tell you the exact date of your heart attack. But the ECG reading is not useless. Quite the opposite.

PART 8: CONNECTING TO THE CMH AXIOMATIC SYSTEM

Why This Matters for the Theory

So far, this guide has treated the Koopman programme primarily as a technical extension of CMH's empirical toolkit. But the connection runs deeper than that. The Koopman–A2 Correspondence conjecture ties the entire programme to the foundational axioms of CMH in a way that makes the empirical test simultaneously a test of the method and a test of the theory.

Axiom A2 — conditional ergodicity — is the axiomatic foundation for everything CMH does comparatively. It is the claim that says: societies operating under similar structural conditions share the same statistical regularities. Without A2, there is no basis for comparing Tunisia to Egypt, or France to Germany, or the Arab Spring to any other historical episode. Every case would be sui generis, incomparable to every other, and the entire scientific programme would collapse into a collection of unconnected case studies.

A2 is, in a sense, the axiom that makes social science possible at all. And until now, it has been largely an assumption — a theoretically motivated postulate that has never been subjected to a direct empirical test within the CMH framework.

The Koopman programme changes that. If the multi-task DMD estimation extracts modes that are genuinely shared across the European panel — if France and Germany and the United Kingdom turn out to oscillate with the same characteristic frequencies, differing only in amplitude and phase — then A2 has passed a concrete empirical test. The assumption of shared dynamics has been confirmed, not just assumed.

If the modes are not shared — if the spectral structures of different countries are incommensurable — then A2 is in trouble, and the entire comparative enterprise requires rethinking.

This is what separates a scientific theory from an ideology. A scientific theory makes predictions that can be wrong, and it takes those predictions seriously. CMH Axiom A2 now has a genuine test in front of it. Whichever way the result goes, the theory will know more about itself than it did before.

What the Axioms Say About the Modes

Beyond A2, the Koopman framework connects to CMH's other axioms in illuminating ways.

Axiom A5 — endogenous indeterminacy — says that complex social systems are genuinely chaotic: their trajectories cannot be predicted indefinitely, regardless of measurement precision. The Koopman framework does not contradict this. It agrees with it, and then asks a different question. A5 governs trajectories. Koopman modes are not trajectories; they are the geometry of the space trajectories inhabit. The axiom that says "you cannot predict the path" says nothing about whether "the valley has a characteristic shape."

Axiom A7 — predictive decay — says that forecast accuracy decays monotonically with horizon. The Koopman framework refines this. It is true for event-level prediction. It is substantially less true for regime-level prediction. And it is not straightforwardly applicable to spectral-level prediction, where the relevant form of uncertainty accumulation is linear phase error rather than exponential trajectory divergence. A7 is not contradicted by the Koopman programme; it is given more precise content.

Axiom A8 — computational falsifiability — is the axiom that demands every CMH claim be expressible in a testable form. The Koopman programme satisfies A8 directly: the conjecture is falsifiable, the test is specified, the criterion for rejection is stated in advance. A mode estimated from training data is predictive if and only if it continues to appear in held-out data and generates accurate phase forecasts. There is no post-hoc rationalisation available. The data will deliver a verdict.

PART 9: THE ROAD AHEAD

What Comes Next

With WP-2026-005 now published, the next step in the research programme is the companion paper — which will become the sixth in the FICSS working paper series — developing the full mathematical treatment of the Koopman-CMH integration. It will contain four major components.

The first is the formal proof of the conditions under which the Koopman–A2 Correspondence conjecture holds within the CMH axiomatic system. This requires techniques from ergodic theory and spectral analysis of stochastic operators — tools that are well-established in mathematics but that have not previously been brought to bear on the specific architecture of the CMH state space.

The second is the derivation of the multi-task Hankel-DMD estimator and its statistical properties in the panel setting. "Hankel-DMD" refers to an extension of the basic DMD algorithm in which the state vector is augmented with temporal lags — yesterday's state appended to today's, and the day before that, and so on. This augmentation increases the effective dimensionality of the observable space, allowing the algorithm to identify modes whose spatial signature is spread across multiple time steps. For social systems, where the response to a structural change may unfold over years rather than days, this extension is essential.

The third is the empirical test on the Western European panel: the actual results of the multi-task Hankel-DMD estimation, including the modes identified, their periods, their growth rates, and their out-of-sample forecast accuracy on held-out data.

The fourth is the implications for the CMH predictive architecture: what the modes, if confirmed, mean for how CMH describes its own predictive range and what kinds of statements it is entitled to make about the long-run dynamics of structurally comparable societies.

The Longer View

Beyond the companion paper, the successful Koopman-CMH integration would open several research directions that are not currently possible.

The most immediate is the extension to non-European historical panels. The Western European case is chosen for the first test because of data availability and relative structural comparability. But the most scientifically interesting tests would involve panels spanning structurally different contexts — post-colonial states in Sub-Saharan Africa, Latin American economies over the twentieth century, East Asian developmental states — where the assumption of shared modes is less obvious and therefore more informative if confirmed.

A second direction is the integration of the Koopman spectral structure with the CMH simulation architecture. If persistent modes are identified and characterised, they can be incorporated into the ensemble simulations that generate level-one predictions, providing a more accurate representation of the system's long-run dynamics as a backdrop against which near-term trajectories are simulated. The Phase One and Phase Three predictions — event-level and spectral-level — would then be complementary components of a single, integrated forecasting system.

A third direction is the use of Koopman mode identification as a diagnostic tool for regime transitions. If a mode that has been stable for decades suddenly changes — its growth rate shifts from neutral to positive, or its spatial signature changes — that change is itself an early warning signal. The spectral structure of the system is revealing that something structural has shifted, even if the event-level dynamics remain apparently calm. This would provide a genuinely new kind of early warning capability, operating at a different timescale from the critical slowing down indicators.

PART 10: LESSONS FOR HOW TO THINK ABOUT HISTORY

Five Things This Changes

Every guide in this series has ended with a set of practical implications — things the reader can take away that change how they think about historical prediction, social science, and the future. This one is no different.

First: the question you ask determines what you can know.

The most important lesson of the prediction hierarchy is not about mathematics. It is about epistemology — the study of how we know what we know. If you ask "where will Country X be in thirty years?" the answer is: we don't know, and nobody does. If you ask "what structural pressures will shape Country X's politics over the next thirty years?" you have a much better chance of a useful answer. If you ask "does Country X show evidence of a recurring demographic-structural oscillation, and if so, where in that oscillation does it currently sit?" you are asking a question that might be answerable even at a fifty-year horizon.

The same reality. Three different questions. Three very different quantities of accessible knowledge.

This is not a philosophical retreat from ambition. It is a strategic advance. By choosing the right questions, you can know things that are genuinely out of reach if you insist on the wrong ones.

Second: cycles are not clockwork.

If the Koopman programme confirms the existence of demographic-structural oscillations in historical data, it would be tempting to use them as a predictive clock — "we are in year 17 of the 25-year cycle, therefore instability will peak in year 20." This is the wrong way to think about it.

Persistent oscillations in chaotic systems are not clockwork. They do not tick with mechanical regularity. Their periods are approximate; the phase accumulates uncertainty over time; they can be disrupted by exogenous shocks of sufficient magnitude. What they provide is not a precise schedule but a structural context. "The demographic-structural dynamics of this society are currently in a phase historically associated with elevated instability risk" is a structural contextualisation, not a prediction. It narrows the space of possibilities; it does not specify one.

Third: the distinction between history repeating and history rhyming.

Mark Twain — or someone who sounded like him — is supposed to have said that history doesn't repeat itself but it often rhymes. The Koopman framework is, in a sense, a mathematical formalisation of that intuition. Rhyming is spectral: two lines rhyme when they share a phonetic pattern, not when they share the same words. Two historical episodes rhyme when they share a structural dynamic — a demographic pressure building toward a release, an institutional legitimacy cycle approaching exhaustion — not when they share specific events or actors.

CMH has always been in the business of identifying structural rhymes rather than narrative repetitions. The Koopman framework gives us the tools to do this with more precision than before.

Fourth: science advances by asking better questions, not just by getting better answers.

One of the striking things about the research programme described in *Beyond the Lyapunov Wall* is that its most important contribution may not be any particular finding, but the set of questions it makes it possible to ask at all. Before the Koopman framework, CMH had no principled way to look for persistent oscillations in historical data. The question "does Western Europe have a twenty-five-year demographic-structural cycle?" was not scientifically askable — not because nobody had thought of it, but because there was no agreed method for answering it rigorously.

The multi-task Hankel-DMD framework makes the question answerable. Whatever the answer turns out to be, the ability to ask it rigorously is a contribution to the science.

Fifth: the honest limit is the beginning, not the end.

The Lyapunov Wall was introduced in the previous guide as CMH's hardest limit. It is worth ending this guide by noting that every new tool described here — slow manifold theory, Critical Slowing Down indicators, Koopman-DMD — was discovered or operationalised *after* researchers fully accepted the hard limits imposed by chaos theory. Not in spite of those limits, but because of them.

Accepting that you cannot see through the wall is the first step toward discovering what you can see around it.

PART 11: QUESTIONS READERS WILL ASK

"Aren't you just finding patterns that aren't there?"

This is the most important methodological objection, and it deserves a serious answer.

Any method powerful enough to find real patterns in complex data is also powerful enough to find spurious ones. This problem — called overfitting, or data dredging — is the methodological plague of data-rich sciences. There are enough historical patterns in a dataset covering ten countries across two centuries that any sufficiently flexible method will find some. The question is whether the patterns found are real or artefactual.

The Koopman-DMD programme addresses this through a strict out-of-sample validation protocol, stated in advance. Modes are identified on a training portion of the data — say, observations up to 1990. A mode is considered predictive only if it continues to appear in held-out data after 1990, and generates phase forecasts that are more accurate than a random-phase baseline. The criterion for rejection is stated explicitly: the predicted oscillation must appear in the held-out data within ±10% of the estimated period, and the phase error must not exceed half a cycle.

This is a demanding standard. It rules out post-hoc pattern-matching. It is exactly the operationalisation of Axiom A8 — computational falsifiability — in the Koopman context. Whether the modes survive this test is not known in advance. If they do not, the programme reports the null result, and the theoretical conclusions are adjusted accordingly.

"Doesn't this assume history repeats? But it doesn't."

History does not repeat in the sense of producing identical events in identical sequences. The French Revolution of 1789 was not a replay of anything that came before, and nothing that has happened since has replicated it. If that is what "history repeats" means, then no, CMH does not assume it.

What CMH assumes — through Axiom A2 — is that *structural dynamics repeat*. The mechanism by which large young adult cohorts, finding limited opportunities for advancement and experiencing deep frustration with existing institutions, generate political instability is the same mechanism whether it is operating in ancient Rome, early modern England, or contemporary Tunisia. The actors change. The specific events change. The cultural context changes profoundly. But the structural dynamics — the way the variables relate to each other, the way pressure builds and releases — follow recognisable patterns.

The Koopman programme is looking for those structural patterns, not for narrative repetitions. It would find a twenty-five-year demographic oscillation even if no two episodes within that oscillation involved the same countries, the same type of government, or the same kind of political event. What the oscillation captures is not "the same thing happening again" but "the same structural dynamic playing out in a new context."

This is the same logic that makes it possible to say that the financial crisis of 2008 "resembled" that of 1929 in structural terms — not because the same banks failed, or the same policies were tried, or the same political figures made the same mistakes, but because the underlying dynamics of credit expansion, asset price inflation, liquidity crisis, and demand collapse followed a recognisable structural pattern.

"What if the modes exist but are too weak to be useful?"

This is a genuinely important practical question. A mode can be statistically real — genuinely present in the data — and still have too small an amplitude, or too imprecise a period, to generate useful predictions. A demographic oscillation that accounts for 3% of the variance in political stability indicators is a real finding but a weak one.

The empirical test will report not just whether modes exist but how much of the variance they explain and how precisely their periods can be estimated. If the results show modes that are statistically significant but practically trivial — too weak to narrow the space of possibilities meaningfully — the programme will report exactly that, together with an assessment of what data volume and quality would be required to identify stronger or more precise modes.

This is the right kind of failure: an honest, informative null result that clarifies what the science can and cannot currently deliver. It is far preferable to overclaiming.

PART 12: THE SHAPE OF THE SCIENCE

What Kind of Thing CMH Is Becoming

We began this guide — and the previous one — with Isaac Asimov's Hari Seldon, the fictional prophet of galactic civilisation. It is worth returning to him one final time, not to compare CMH to psychohistory but to clarify the relationship between ambition and honesty in science.

Seldon's psychohistory was built on a single, breathtaking premise: that the future of human civilisation could be calculated, in its precise trajectory, across millennia. This premise made the Foundation series one of the most beloved works of science fiction in history. It also made psychohistory impossible — not as a fictional device, which it remains magnificent, but as a scientific programme.

The predecessor of CMH had the same breathtaking premise. The Lyapunov Wall demolished it.

What emerged from the demolition was not defeat but clarification. The science that CMH is becoming is defined not by what it can predict but by the precision with which it characterises what it cannot predict — and then methodically works on what it can.

It can tell you, with some confidence, what structural pressures are building in a society right now. It can tell you what kind of structural dynamics have historically been associated with elevated instability risk. It can detect early warning signals that a system is approaching a tipping point. It can, potentially, tell you where a society sits in a recurring structural oscillation — which phase of the tide it is currently in.

None of this is prophecy. All of it is useful.

The science is one that knows, with increasing precision, the shape of its own ignorance. That shape is not a defeat. It is a map. And maps — even maps that show where you cannot go — are enormously valuable when you are trying to navigate.

The Last Analogy

There is one more analogy I want to leave with you, because it captures something that none of the previous ones quite reaches.

Astronomers who study pulsars — rapidly rotating neutron stars — face a curious situation. They cannot see a pulsar directly. It is too small, too distant, and too dim. But they can detect the beam of electromagnetic radiation it emits, which sweeps across the Earth with each rotation like the beam of a lighthouse.

By timing the sweeps precisely, astronomers can characterise the pulsar's rotation rate with extraordinary accuracy — sometimes to fifteen or more significant figures. They can predict, centuries in advance, exactly when the next pulse will arrive. Not because they can see the pulsar, but because they have identified the spectral structure of its dynamics: the fundamental oscillation that organises all the observable behaviour.

Now suppose the pulsar were embedded in a turbulent nebula — a cloud of gas and plasma that scattered, absorbed, and randomised the beam on a pulse-by-pulse basis. Each individual pulse would be unpredictable; the turbulence would destroy the signal. But the underlying oscillation — the rotation rate of the pulsar — would still be there, still driving the system, still detectable in principle by sufficiently clever analysis of the statistics of what gets through the nebula.

This is the situation CMH is in. The socio-political systems we study are the turbulent nebula. The chaotic surface dynamics — the specific events, the specific triggers, the specific trajectories — scatter and randomise in ways that defeat event-level prediction beyond the Lyapunov horizon. But if there are pulsars in there — persistent structural oscillations that drive the system's dynamics at a deep level — then they are detectable, in principle, by exactly the methods this guide has described.

Beyond the Lyapunov Wall is the proposal to look. The companion paper will report what we find.

This guide is the fifth in the FICSS popular science series on Computational Macrohistory. The previous guides explored the axiomatic foundations of the framework, the Lyapunov Wall and its implications, and the master equation that describes how societies change. This guide explored the new research programme introduced in WP-2026-005, available at Zenodo (DOI: 10.5281/zenodo.20111272), and previewed the companion paper that will develop its mathematical foundations in full.

Galen Fontaise is the founder and director of the Fontaise Institute of Computational Social Science (FICSS) in Lugano, Switzerland. The full CMH working paper series is available at ficss.ch and on Zenodo.