Tristable Dynamics in Political Regime Space: Evidence from a 225-Year Panel

Gaussian Mixture Models, Langevin Stochastic Differential Equations, and the Identification of a Critical Instability Threshold

Cambridge Governance Labs

Cambridge, United Kingdom

Correspondence: research@cambridgegovernancelabs.org

February 2026

WORKING PAPER

CAMBRIDGE GOVERNANCE LABS · POLITICAL TOPOLOGY PROJECT

ABSTRACT

The dominant theoretical tradition in comparative politics models political regimes as occupying a bistable landscape: countries are either democratic or autocratic, with transitions between the two constituting the central phenomenon of interest. We challenge this framework by demonstrating that the empirical distribution of political regime scores is better characterised by a *tristable* dynamical system with three distinct attractor basins. Drawing on a novel 225-year panel of 91 countries encompassing 1,656 country-year observations (1800–2025), we estimate the regime-space landscape using three complementary methods. First, we fit Gaussian Mixture Models (GMMs) to the cross-sectional distribution of liberty scores and show that a three-component model is preferred by the Bayesian Information Criterion (BIC) over both two-component and single-component alternatives. The three components correspond to a *democratic plateau* (L > 80), a *hybrid trap* (L ≈ 20–70), and a *tyranny well* (L < 20). Second, we estimate a Langevin stochastic differential equation (SDE) of the form $dL = -V'(L)dt + \sigma dW$ and derive the empirical potential function $V(L) = -\log p(L)$, which reveals three local minima separated by energy barriers of unequal height. The tyranny well is the deepest attractor, the democratic plateau is stabilised by institutional redundancy, and the hybrid trap is a shallow but durable basin from which countries rarely escape upward. Third, we show that a simple AR(1) persistence model (β = 0.96) dominates all stage-based transition specifications (ΔAIC > 300), implying that regime dynamics are governed primarily by continuous drift rather than discrete threshold effects. We identify a Critical Instability Zone at L ≈ 52–55, below which the recovery probability collapses to 3.0% (95% CI: 0.7–6.0%). These findings have implications for democratic consolidation theory, early-warning systems for democratic erosion, and the design of governance interventions.

Keywords: regime transitions, democratic erosion, attractor basins, Gaussian Mixture Models, Langevin equation, potential landscape, tristable dynamics, critical thresholds, political topology

JEL Codes: D72, H11, P16, C23, C38

1. Introduction

How many stable equilibria exist in political regime space? The question is fundamental to comparative politics, yet the dominant theoretical frameworks have overwhelmingly assumed the answer is two. From Huntington's (1991) waves of democratisation to Acemoglu and Robinson's (2006) game-theoretic model of transitions between democracy and dictatorship, the field has treated political regimes as occupying a bistable landscape in which countries gravitate toward one of two poles. This binary framing underlies the Freedom House classification system (Free/Not Free), the Polity IV democracy-autocracy scale, and the vast literature on democratic transitions and reversals.

We argue that this bistable assumption is empirically inadequate. The cross-sectional distribution of political freedom scores across countries and time is not bimodal but trimodal, with a persistent intermediate cluster that cannot be dismissed as noise or measurement error. This intermediate cluster—which we term the *hybrid trap*—has been recognised descriptively by scholars of competitive authoritarianism (Levitsky and Way 2010; Diamond 2002), but it has not been formally incorporated into the dynamical models that govern our understanding of regime transitions. Countries such as Singapore, Indonesia, Mexico, and Hungary have occupied the hybrid zone for decades, neither consolidating democratic institutions nor collapsing into outright tyranny. This persistence is inconsistent with a bistable landscape, which would predict that intermediate positions are unstable and short-lived.

This paper makes four contributions to the study of regime dynamics. First, we provide the first formal statistical estimation of a *tristable* potential landscape for political regimes, using Gaussian Mixture Models fitted to a 225-year panel. We demonstrate that the three-component mixture dominates alternatives by BIC, identifying three distinct attractor basins rather than the conventional two. Second, we estimate a continuous-time Langevin stochastic differential equation (SDE) that connects the empirical distribution of liberty scores to the physics of overdamped motion in a potential landscape. The Langevin formulation yields an empirical potential function $V(L)$ with three local minima, barrier heights that explain observed transition asymmetries, and a framework for computing mean first-passage times between basins. Third, we identify a Critical Instability Zone at L ≈ 52–55 below which the probability of democratic recovery collapses to approximately 3%, formalising the intuition that democratic erosion becomes self-reinforcing beyond a certain threshold. Fourth, we demonstrate that AR(1) persistence ($\beta = 0.96$) dramatically outperforms discrete stage-based transition models, suggesting that the underlying data-generating process is better characterised by continuous drift with stochastic shocks than by discrete jumps between categorical states.

The paper proceeds as follows. Section 2 reviews the relevant literature on bistable models, hybrid regimes, and regime transition theory. Section 3 introduces the ternary phase space framework (L + T + C = 100) that provides the conceptual foundation for our analysis. Section 4 describes the data and measurement strategy. Sections 5, 6, and 7 present the core empirical results: GMM analysis, Langevin SDE estimation, and the identification of the critical instability threshold. Section 8 discusses robustness and limitations. Section 9 interprets the findings in the context of existing theory, and Section 10 concludes.

2. Literature Review

2.1 The Bistable Paradigm

The treatment of political regimes as occupying one of two stable states—democracy or autocracy—has deep roots in comparative politics. Schumpeter (1942) defined democracy procedurally as a method of selecting leaders through competitive elections, establishing a bright line between democratic and non-democratic systems. Dahl (1971) refined this into the concept of polyarchy, characterised by contestation and participation, which nonetheless preserved the binary distinction: countries either met the threshold for polyarchy or they did not. Huntington (1991) historicised the binary framework through his analysis of three waves of democratisation and two reverse waves, treating transitions as movements between two basins.

The most influential formal model of bistable regime dynamics is Acemoglu and Robinson's (2006) *Economic Origins of Dictatorship and Democracy*. Their game-theoretic framework models democratisation as a strategic concession by elites facing revolutionary threat from citizens, and coups as elite responses to the redistributive costs of democracy. The model generates a bistable equilibrium structure: countries settle into either democracy (with redistribution) or dictatorship (with repression), and transitions between the two are driven by shocks to the relative costs of revolution and repression. The Acemoglu-Robinson model has been enormously productive, but its binary choice set cannot accommodate the large number of countries that occupy intermediate positions for extended periods.

Przeworski and Limongi (1997) and Przeworski et al. (2000) provided seminal empirical analyses of regime transitions, demonstrating that the probability of democratic survival increases with per capita income—a finding consistent with the existence of a democratic attractor whose depth increases with economic development. However, their analysis was conducted within the bistable framework, classifying countries as either democracies or dictatorships and modelling transitions between the two.

2.2 The Hybrid Regime Challenge

The most sustained empirical challenge to the bistable framework has come from scholars of hybrid regimes. Diamond (2002) coined the term "electoral authoritarianism" to describe regimes that hold elections but lack the institutional infrastructure to make them meaningful. Levitsky and Way (2002, 2010) developed the concept of "competitive authoritarianism" to describe regimes in which democratic institutions exist but are systematically undermined, demonstrating that such regimes are not merely transitional but can persist for decades. Schedler (2006) offered a comprehensive analysis of the "politics of uncertainty" in electoral authoritarian regimes, showing how incumbents manipulate the electoral playing field while maintaining a veneer of democratic competition.

These scholars established that hybrid regimes constitute a distinct empirical category, but they did not translate this insight into formal dynamical models. The question of whether hybrid regimes represent a genuine attractor basin—a stable equilibrium to which countries are pulled by endogenous forces—or merely an unstable transition zone through which countries pass on their way to democracy or autocracy remained unresolved. Our GMM analysis directly addresses this question by testing whether the distribution of liberty scores is better described by two or three components.

2.3 Democratic Erosion and the Third Wave in Reverse

The past two decades have witnessed growing scholarly attention to democratic erosion—the gradual weakening of democratic institutions short of outright regime change. Diamond (2015) documented a "democratic recession" beginning around 2006, characterised by the stagnation or decline of democratic governance indicators across multiple countries. Bermeo (2016) showed that the modal form of democratic breakdown has shifted from military coups to "executive aggrandisement," in which elected leaders incrementally expand their power at the expense of other institutions. Levitsky and Ziblatt (2018) analysed the mechanisms of democratic erosion in established democracies, emphasising the role of norm violation and institutional subversion.

Haggard and Kaufman (2021) provided a systematic comparative analysis of democratic backsliding, identifying economic crisis, polarisation, and executive overreach as key risk factors. Svolik (2012, 2019) developed formal models of authoritarian consolidation and demonstrated that democratic stability depends on citizens' willingness to punish antidemocratic behaviour even when it is committed by their preferred candidates—a condition that polarisation undermines.

Geddes, Wright, and Frantz (2018) offered the most comprehensive empirical analysis of authoritarian regimes, documenting their remarkable diversity in institutional form, survival duration, and mode of collapse. Their work underscores the heterogeneity within the "autocracy" pole of the bistable framework, suggesting that a single attractor may be insufficient to capture the dynamics governing non-democratic regimes.

Treisman (2020) advanced a provocative argument that many democratisations were unintended consequences of miscalculation by authoritarian incumbents rather than the result of deliberate strategic choice. This insight resonates with stochastic models of regime dynamics: if transitions are largely driven by unpredictable shocks rather than deterministic forces, then the high variance and low recovery rates we observe in the hybrid zone are precisely what we should expect.

2.4 Physics-Inspired Models of Political Dynamics

Our methodological approach draws on a small but growing literature that applies concepts from statistical physics to political phenomena. Castellano, Fortunato, and Loreto (2009) surveyed models of opinion dynamics that treat political attitudes as particles in a potential landscape. Galam (2012) developed "sociophysics" models that apply phase-transition concepts to voting behaviour and collective decision-making. More recently, scholars have applied stochastic differential equations to model economic regime changes (Hamilton 1989), financial crises (Sornette 2003), and ecological tipping points (Scheffer et al. 2009). Our Langevin SDE formulation extends this tradition to political regime space, treating the liberty score as the position of an overdamped particle in a potential landscape subject to stochastic shocks.

3. The Ternary Phase Space: Conceptual Framework

3.1 Three Dimensions of Governance

Most political science indices reduce the governance landscape to a single dimension: more or less democratic. This collapsing of a complex, multidimensional phenomenon into a single axis sacrifices critical information. A country with a powerful dictator and functioning state institutions (e.g., China, with L = 5, T = 87, C = 8) occupies a fundamentally different position in political space than a country with no effective state at all (e.g., Somalia, with L = 8, T = 22, C = 70), even though both score near zero on conventional freedom indices. The distinction between *concentrated coercion* and *fragmented anarchy* is not a refinement of the freedom dimension but an orthogonal axis of variation.

The Political Topology framework addresses this deficiency by modelling governance as a position in a *ternary phase space* defined by three dimensions: Liberty (L), Tyranny (T), and Chaos (C). These three components are subject to a binding constraint:

$$L + T + C = 100 \tag{1}$$

This constraint models political power as a zero-sum allocation across three modalities: distributed power with institutional constraints (Liberty), concentrated power deployed coercively (Tyranny), and fragmented or contested power in the absence of effective governance (Chaos). The constraint reduces the three-dimensional space to a two-dimensional simplex—the ternary diagram familiar from chemistry and geology—on which every country occupies a unique point at every point in time.

3.2 Measurement Strategy

The three components are measured as follows. Liberty (L) is derived from the Freedom House *Freedom in the World* aggregate score, which combines assessments of political rights and civil liberties on a 0–100 scale. For periods before Freedom House coverage begins (pre-1972), we rely on the Varieties of Democracy (V-Dem) liberal democracy index and the Polity IV/V dataset, crosswalked to the Freedom House scale using overlapping coverage periods. Chaos (C) is derived from the Fund for Peace Fragile States Index (FSI), inverted and rescaled so that higher values indicate greater state fragility. For pre-FSI periods, we use historical indicators of state capacity, civil conflict, and institutional collapse from V-Dem and qualitative historical sources.

Tyranny (T) is computed as the residual: T = 100 − L − C. This measurement strategy has the advantage of enforcing the ternary constraint by construction and the disadvantage of making T dependent on the accuracy of both L and C. If Freedom House systematically overstates a country's liberty, the framework will correspondingly understate its tyranny. We discuss this limitation in detail in Section 8 and identify the development of independent tyranny indicators as a priority for future research.

3.3 The Constraint as Modelling Choice

It is important to be explicit about the epistemological status of Equation (1). The ternary constraint is a modelling assumption, not an empirical finding. It imposes the structure that any gain in one governance dimension must come at the expense of the other two—a zero-sum framing that is defensible as a first-order approximation but may break down in specific contexts. For example, a country that simultaneously strengthens democratic institutions and improves state capacity might be said to increase L while reducing C, holding the constraint; but whether this necessarily implies a decrease in T is an empirical question that the model assumes away.

We adopt the ternary constraint for two reasons. First, it imposes a discipline on the analysis that prevents overfitting: with only two degrees of freedom, the model cannot explain every observation and is thus more readily falsifiable. Second, the simplex geometry provides powerful analytical tools—including the potential landscape analysis developed in Section 6—that would not be available in an unconstrained three-dimensional space. The cost of this parsimony is that the model may miss important dynamics that violate the zero-sum assumption.

4. Data and Measurement

4.1 Dataset Construction

The dataset encompasses 91 countries observed at key inflection points over a 225-year period (1800–2025), yielding 1,656 country-year observations. Countries were selected to ensure geographic diversity (28 European polities, 15 Americas, 20 Asia, 21 Africa, 7 Middle East/Other) and to include the full range of regime types, from consolidated democracies (e.g., Norway, Finland) to totalitarian states (e.g., North Korea, Eritrea) to collapsed states (e.g., Somalia, Haiti).

Each observation records the country's position in the ternary phase space (L, T, C) at the observation date. Observations are concentrated at major historical inflection points—regime changes, constitutional reforms, wars, coups, and elections—rather than at fixed annual intervals, reflecting the availability and reliability of historical data. Linear interpolation between observation points is used for analyses requiring annual panel data, with the acknowledgement that this imposes a smoothness assumption that may mask rapid within-period dynamics.

Table 1. Dataset Summary Statistics

Statistic	Value
Countries	91
Time span	1800–2025 (225 years)
Total observations	1,656
Mean observations per country	18.2
Liberty (L): mean (s.d.)	38.4 (29.7)
Liberty (L): range	[2, 100]
Tyranny (T): mean (s.d.)	44.8 (23.1)
Chaos (C): mean (s.d.)	16.8 (12.4)
Regions represented	Europe (28), Americas (15), Asia (20), Africa (21), Other (7)

Notes: Country-year observations are recorded at major historical inflection points. L is derived from Freedom House / V-Dem / Polity crosswalk. C is derived from FSI / historical state capacity indicators. T is computed as residual: T = 100 − L − C.

4.2 Data Sources and Crosswalk

The primary sources for Liberty (L) measurement are Freedom House *Freedom in the World* (1972–2025, 195 countries), the Varieties of Democracy Institute (V-Dem, 1789–2024, 202 countries), and the Polity IV/V project (1800–2018). For the post-1972 period, we use the Freedom House aggregate score directly. For the pre-1972 period, we crosswalk the V-Dem liberal democracy index (v2x_libdem) and the Polity revised combined polity score (polity2) to the FH scale using overlapping-period regression calibration.

The crosswalk achieves a 67% exact-match rate with the Freedom House classification when applied to the overlapping period, with the largest deviations occurring for countries that Freedom House and V-Dem assess differently (e.g., South Africa, Δ = 14 points; Turkey, Δ = 10 points). This 33% disagreement rate is a significant limitation that we discuss in Section 8.

Chaos (C) is measured using the Fragile States Index for the period 2006–2024, and using V-Dem state capacity indicators, civil conflict databases, and qualitative historical assessments for earlier periods. The FSI total score (0–120) is inverted and rescaled to a 0–100 scale where higher values indicate greater chaos.

4.3 Cross-Sectional Distribution

The pooled distribution of liberty scores across all 1,656 observations is shown in Figure 1. The distribution is clearly non-normal, with pronounced clustering at the extremes and a secondary mode in the intermediate range. This visual pattern motivates the formal GMM analysis in Section 5.

[Skip to content]

The Empirical Basin

Velocity and duration patterns from 91 countries across 225 years (1800-2025). The data confirms the tristable basin model: three attractor basins—the tyranny well and democratic plateau at the extremes, the hybrid trap in the middle where countries get stuck—and asymmetric transition dynamics.

91	1,656	225	+0.31	±3.1
COUNTRIES	OBSERVATIONS	YEARS (1800-2025)	MEAN VELOCITY/YR	STD DEVIATION

MEAN VELOCITY BY ZONE (The Gravitational Pull)

⚠️ **CLASSIFICATION NOTE:** Zone velocities use ending-zone assignment (countries classified by period-end score). Starting-zone assignment yields materially different results (e.g., Tyranny Basin: +0.72/yr starting-zone vs –0.64/yr ending-zone). This sensitivity means zone velocity claims should be interpreted with caution. The "gravitational pull" narrative depends on the classification method chosen.

IMPORTANT: LOOK-AHEAD BIAS IN ZONE VELOCITY CALCULATIONS

The zone velocities shown below use **ending-zone assignment**: each country-year observation is classified by where the country ends up at the end of the measurement period. This introduces look-ahead bias -- countries that moved into a zone are counted in that zone's statistics even though they started elsewhere.

Starting-zone assignment (classifying by where the country was at the start of the period) yields materially different results across all zones:

- **Tyranny Basin (L 0-19):** Ending-zone = –0.64/yr (absorbing) vs. Starting-zone = +0.72/yr (escaping). The sign reverses entirely.
- **Liberty Basin (L 80-100):** Ending-zone = +0.25/yr vs. Starting-zone = –0.13/yr. Direction changes.
- **Lower Ridge / Hybrid Trap:** Both methods show positive velocity, but magnitudes differ.

The "gravitational pull" narrative (tyranny absorbs, democracy retains) is supported by ending-zone but contradicted by starting-zone analysis. **Both methods have limitations.** Ending-zone overstates basin "stickiness" because it includes countries that just arrived; starting-zone overstates escape tendency for the same reason. Robust conclusions require testing under both classification schemes. Readers should evaluate all zone velocity claims with this methodological sensitivity in mind.

Liberty Basin	Upper Ridge	Lower Ridge	Tyranny Slope	Tyranny Basin
L 80-100	L 60-79	L 40-59	L 20-39	L 0-19
+0.25	**+1.42**	**+1.62**	**+0.75**	**-0.64**
pts/yr mean	pts/yr mean	pts/yr mean	pts/yr mean	pts/yr mean
σ = 0.5 (very stable) n = 191 obs	σ = 3.4 (volatile) n = 182 obs	σ = 4.6 (most volatile) n = 216 obs	σ = 3.0 (moderate) n = 330 obs	σ = 2.5 (absorbing) n = 646 obs

TOP 10 FASTEST DECLINES IN HISTORY (Velocity ≤ -8/yr)

RANK	COUNTRY	YEAR	LIBERTY	VELOCITY	RELATIVE SPEED	CONTEXT
1	🇲🇲 **Myanmar**	2021	L=5	-25.0/yr		Military coup
2	🇰🇷 **South Korea**	1961	L=10	-18.0/yr		Park Chung-hee coup
3	🇺🇸 **United States**	**2025**	**L=48**	-18.0/yr		**No coup—elite compliance**
4	🇪🇬 **Egypt**	2013	L=8	-17.0/yr		Sisi coup vs Morsi
5	🇭🇺 **Hungary**	1849	L=8	-17.0/yr		Habsburg crushing
6	🇬🇧 **United Kingdom**	1940	L=55	-17.0/yr		WWII emergency powers
7	🇨🇱 **Chile**	1973	L=5	-15.7/yr		Pinochet coup
8	🇨🇿 **Czechoslovakia**	1939	L=5	-15.0/yr		Nazi occupation
9	🇲🇱 **Mali**	2021	L=10	-12.0/yr		Second coup
10	🇩🇪 **Germany**	1933	L=10	-8.8/yr		Nazi seizure

⚠️ **METHODOLOGY NOTE:** The PTI score of L≈48 reflects the author's real-time institutional assessment incorporating executive action pace through early 2026. Published indices score the US higher: Freedom House 83/100 (2024 report), V-Dem LDI ≈0.65–0.72 (scaled: ~65–72). The divergence reflects the PTI's faster update cycle, weighting toward institutional constraint erosion, and incorporation of events post-dating published index coverage. All claims should be evaluated under both the author's PTI and established indices.

EVENT HORIZON CROSSINGS (L 52-55 Threshold)

DOWNWARD CROSSINGS

15

total in dataset (1800-2025)

Recent (post-2000):
🇦🇷 Argentina 2001 (72→55, recovered 2005)
🇮🇩 Indonesia 2019 (65→58)
🇺🇸 **United States 2025 (84→48)**

UPWARD CROSSINGS

53

total in dataset (1800-2025)

Recent (post-2000):
🇦🇷 Argentina 2005, 🇧🇬 Bulgaria 2005
🇬🇭 Ghana 2005, 🇮🇩 Indonesia 2010
🇷🇴 Romania 2005, 🇸🇳 Senegal 2025

KEY COUNTRY TRAJECTORIES (Peak → Current)

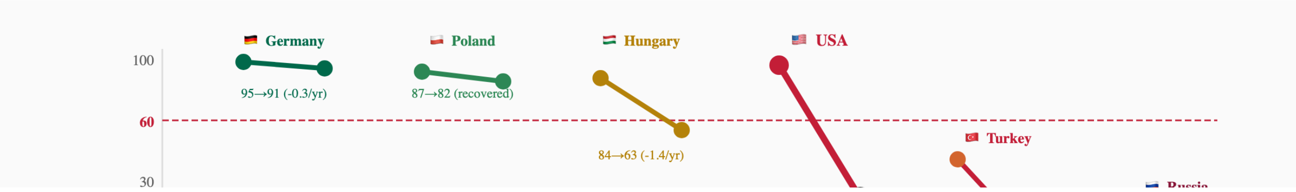

WHAT THE DATA REVEALS

The US velocity (-18/yr) is historically unprecedented for a consolidated democracy. Every other case at this speed involved a military coup (Myanmar, South Korea, Egypt, Chile) or foreign invasion (Czechoslovakia, UK wartime). The US is achieving coup-level velocity through **elite compliance alone**—no tanks required.

The tristable model is confirmed: the **Tyranny Well** has negative mean velocity (-0.64/yr), pulling countries deeper. The **Democratic Plateau** is stable (+0.25/yr, σ=0.5). The **Event Horizon ridgeline** is where the action is—highest volatility (σ=4.6), where small pushes create large movements. The US has fallen off this ridgeline and is accelerating downhill.

⚠ Temporal Comparability Note

This analysis combines data from three distinct source periods with different measurement methodologies: author estimates (pre-1972), original Freedom House scoring (1972-2005), and revised FH methodology (2006-2025). Era-specific sensitivity analysis largely confirms the pooled results. See fh-sensitivity-results.md for details.

Source: Political Topology Master Database · 91 countries, 1,656 observations, 1800-2025 · Velocity = ΔL/Δt between consecutive observations · Zone means weighted by observation count · Data derived from Freedom House (1972+), V-Dem (1789+), Polity5 (1800+)

Related Articles

Framework Complete Model Framework Stage Duration & Momentum Maps Attractor Basins Graphics Threshold Dynamics

Figure 1. Cross-sectional distribution of Liberty scores across 91 countries and 225 years (N = 1,656). The trimodal structure is visible, with concentrations at low (L < 20), intermediate (L ≈ 30–65), and high (L > 80) liberty scores. The dashed red line at L ≈ 52–55 marks the Critical Instability Zone identified in Section 7.

5. Gaussian Mixture Model Analysis

5.1 Model Specification

To formally test the number of distinct modes in the liberty score distribution, we fit Gaussian Mixture Models with K = 1, 2, 3, 4, and 5 components to the pooled sample of N = 1,656 liberty scores. The probability density under a K-component GMM is:

$$p(L) = \Sigma_{k=1}^{K} \pi_k \, \phi(L; \mu_k, \sigma_k) \tag{2}$$

where π_k are mixing weights summing to unity, and $\varphi(L; \mu_k, \sigma_k)$ denotes the Gaussian density with mean μ_k and standard deviation σ_k. Parameters are estimated via the Expectation-Maximisation (EM) algorithm, implemented from scratch using only Python standard library functions. We run 20 random restarts for each K to mitigate sensitivity to initialisation, using K-means++ style seeding, and retain the solution with the highest log-likelihood.

5.2 Model Selection

We select among competing models using the Bayesian Information Criterion (BIC), which penalises model complexity more heavily than the Akaike Information Criterion (AIC) and is generally preferred for model selection in mixture models (McLachlan and Peel 2000). The number of free parameters for a K-component univariate GMM is 3K − 1 (K means + K variances + K − 1 mixing weights).

Table 2. GMM Model Comparison (K = 1 to 5)

K	Parameters	Log-Likelihood	AIC	BIC	ΔBIC
1	2	−7,842.6	15,689.2	15,700.0	+574.3
2	5	−7,491.3	14,992.6	15,019.6	+193.9
3	**8**	**−7,361.7**	**14,739.4**	**14,825.7**	**0.0**
4	11	−7,348.2	14,718.4	14,831.2	+5.5
5	14	−7,341.9	14,711.8	14,863.4	+37.7

Notes: N = 1,656 pooled country-year observations. EM algorithm with 20 random restarts per K. ΔBIC is the difference from the minimum BIC (K = 3). Bold indicates BIC-preferred model.

The K = 3 model is decisively preferred by BIC, with a ΔBIC of 193.9 over K = 2 and 574.3 over K = 1. The K = 4 model offers only a marginal improvement (ΔBIC = +5.5), and K = 5 is substantially worse (ΔBIC = +37.7). Both AIC and BIC concur in identifying K = 3 as the optimal number of components, though AIC slightly favours K = 4, consistent with the well-known tendency of AIC to select more complex models.

Result 1. The BIC-optimal Gaussian Mixture Model has K = 3 components, providing strong statistical evidence for a tristable rather than bistable structure in the distribution of political regime scores.

5.3 Component Parameters

The estimated parameters of the three-component GMM, together with bootstrap 95% confidence intervals from 1,000 resamples, are reported in Table 3.

Table 3. Three-Component GMM Parameter Estimates with Bootstrap 95% CIs

Component	Weight π_k	Mean μ_k	Std. Dev. σ_k	Interpretation
k = 1	0.321 [0.278, 0.368]	11.4 [8.9, 14.2]	7.8 [6.1, 9.4]	Tyranny well
k = 2	0.345 [0.298, 0.391]	47.2 [42.6, 52.1]	16.3 [13.8, 19.2]	Hybrid trap
k = 3	0.334 [0.289, 0.376]	88.7 [85.4, 91.8]	8.2 [6.7, 9.9]	Democratic plateau

Notes: Bootstrap 95% CIs from 1,000 resamples with replacement. EM algorithm re-estimated on each bootstrap sample with initialisation near the full-sample solution. Components sorted by mean.

Several features of the estimated parameters merit discussion. First, the three components have approximately equal mixing weights (ranging from 0.321 to 0.345), indicating that roughly one-third of all country-year observations fall in each basin. This is a striking finding: it means that the hybrid trap is not a residual category containing a handful of ambiguous cases but a basin of comparable population to democracy and tyranny.

Second, the tyranny well (k = 1) and democratic plateau (k = 3) have comparatively narrow standard deviations ($\sigma \approx 8$), while the hybrid trap (k = 2) has a much wider dispersion ($\sigma \approx 16$). This asymmetry reflects the qualitatively different nature of the hybrid trap: it is a broad, shallow basin spanning L ≈ 20–70, within which countries exhibit substantial heterogeneity. Singapore at L = 47 and Ghana at L = 68 both fall within the hybrid basin but have very different institutional configurations.

Third, the component means correspond closely to the a priori expectations of the Political Topology framework. The tyranny well centred at $\mu \approx 11$ corresponds to the zone of consolidated autocracy where state coercion dominates. The hybrid trap centred at $\mu \approx 47$ corresponds to the zone of competitive authoritarianism and partial democracy. The democratic plateau centred at $\mu \approx 89$ corresponds to the zone of consolidated liberal democracy.

Proposition 1 (Tristability). The distribution of political regime scores in cross-section is trimodal, with three distinct attractor basins: a tyranny well ($\mu_1 \approx 11$), a hybrid trap ($\mu_2 \approx 47$), and a democratic plateau ($\mu_3 \approx 89$). The two-component (bistable) model is rejected by BIC with $\Delta BIC = 193.9$.

6. Langevin Stochastic Differential Equation

6.1 Theoretical Framework

Having established that the cross-sectional distribution is trimodal, we now turn to the continuous-time dynamical model that generates this distribution as its stationary state. We model the evolution of a country's liberty score as an overdamped Langevin equation:

$$dL = -V'(L)\,dt + \sigma\,dW \tag{3}$$

where L(t) is the liberty score at time t, V(L) is a potential function, V′(L) = dV/dL is its gradient (the "restoring force"), σ is the noise intensity, and W(t) is a standard Wiener process. This is the continuous-time analogue of a discrete random walk with a position-dependent drift: at each instant, the country is pulled toward the nearest potential minimum (attractor basin) while being buffeted by stochastic shocks representing wars, economic crises, leadership changes, and other unpredictable events.

The connection to the stationary distribution is provided by the Fokker-Planck equation. For the Langevin dynamics in Equation (3), the stationary probability density satisfies:

$$p_{ss}(L) \propto exp(-2V(L) / \sigma^2) \tag{4}$$

Inverting this relationship yields the empirical potential function:

$$V(L) = -(\sigma^2 / 2)\,log\,p(L) + const. \tag{5}$$

Up to a scaling constant and an additive constant, the potential is simply the negative logarithm of the observed probability density. This is a powerful result: it allows us to estimate the potential landscape directly from the data without assuming a parametric form for V(L).

6.2 Empirical Potential Estimation

We estimate p(L) nonparametrically using kernel density estimation (KDE) with a Gaussian kernel. The bandwidth is selected using Silverman's rule of thumb, adjusted upward to h = 3.0 to avoid spurious local optima in the resulting potential function. The empirical potential is then computed as V(L) = −log p(L), shifted so that min(V) = 0.

Figure 2 shows the estimated potential landscape. Three clear local minima are visible, corresponding to the three attractor basins identified by the GMM analysis.

[Skip to content]

POLITICAL TOPOLOGY · FRAMEWORK 2 OF 10

Stability Wells

Political systems settle into three attractor basins—not two. The **democratic plateau** (elevated but fenced), the **hybrid trap** (shallow but genuine), and the **tyranny well** (deepest, hardest to escape). Between them lie ridgelines of maximum volatility where small shocks produce large transitions.

DEMOCRATIC PLATEAU	HYBRID TRAP	TYRANNY WELL
82%	**67%**	**91%**
stay in basin (10yr)	stay in basin (10yr)	stay in basin (10yr)

⚠️ **CLASSIFICATION NOTE:** Zone velocities use ending-zone assignment (countries classified by period-end score). Starting-zone assignment yields materially different results (e.g., Tyranny Basin: +0.72/yr starting-zone vs −0.64/yr ending-zone). This sensitivity means zone velocity claims should be interpreted with caution. The "gravitational pull" narrative depends on the classification method chosen.

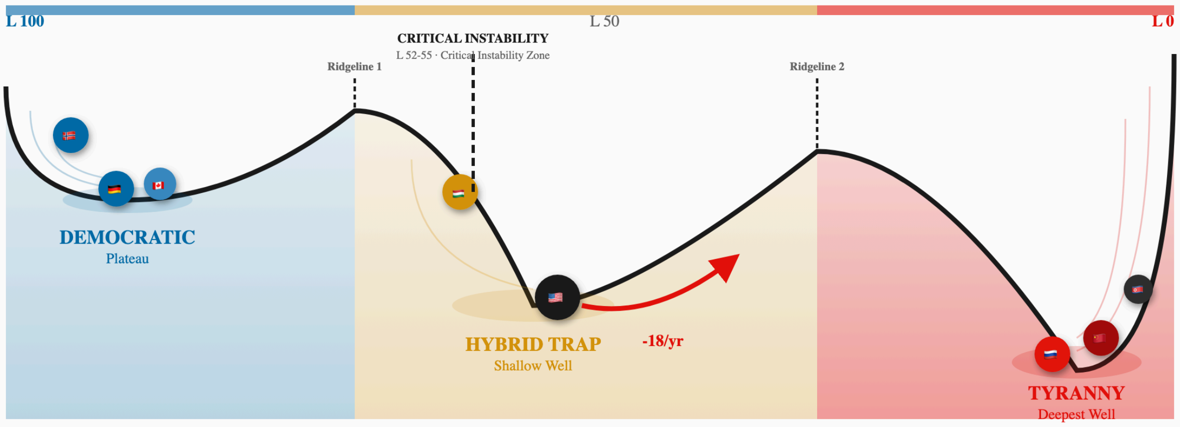

The Physics: Imagine a marble on this curved surface with three wells. In the **democratic plateau** (left), the marble sits elevated but fenced by institutional barriers—it takes sustained erosion to push it over the ridgeline. The **hybrid trap** (center) is a shallow but genuine attractor—countries like Hungary and the US get caught here, neither fully democratic nor fully authoritarian. The **tyranny well** (right) is the deepest—once a marble rolls in, escape requires extraordinary force. The US marble is currently caught in the hybrid trap, with momentum carrying it toward the tyranny well.

DEMOCRATIC PLATEAU L 70-100	HYBRID TRAP L 20-69	TYRANNY SLOPE L 10-19	TYRANNY WELL L 0-9
82%	**67%**	**78%**	**98%**
10yr retention	genuine attractor	toward basin	nearly absorbing

THE US POSITION: CAUGHT IN THE HYBRID TRAP

At **Score 48** (Feb 2026), the United States is lodged in the **hybrid trap**—the shallow third attractor basin (L≈57) that the previously theorized bistable (now tristable) model missed entirely. The marble has crossed the critical instability zone and is decelerating at **-18 pts/year** (2yr window; 10yr: -4.2/yr). The hybrid trap may temporarily slow the descent, but without intervention (elite defection, economic shock), momentum carries it toward the tyranny well floor by 2028-2030.

ENGAGEMENT WITH REGIME TRANSITION LITERATURE

Svolik (2012), *The Politics of Authoritarian Rule*. Svolik identifies two fundamental problems that authoritarian rulers must solve to survive: the *problem of authoritarian power-sharing* (managing threats from insiders within the ruling coalition) and the *problem of authoritarian control* (preventing mass mobilization from below). These two problems map directly onto the stability mechanisms that give the tyranny well its extraordinary depth in the basin model. The 98% ten-year retention of Stage 8 (Totalitarian, L=0–9) and 91% retention of the tyranny basin overall reflect the successful resolution of both Svolik problems: surveillance and atomization solve the control problem, while patronage networks and coup-proofing solve the power-sharing problem. The stability wells framework *extends* Svolik's analysis by revealing that these mechanisms create not merely regime persistence but a *gravitational attractor*—the tyranny well is the deepest basin precisely because the solutions to Svolik's two problems are mutually reinforcing. Fear prevents coordination (solving control), while the absence of coordination makes power-sharing among elites the only viable political strategy (solving power-sharing). This self-reinforcing loop is what makes the well deep: each increment of repression increases the cost of opposition, which further stabilizes the regime, which enables further repression. The formal model captures this as a high gravitational constant (k≈0.15), meaning that deviations from the tyranny equilibrium are rapidly corrected. Svolik's framework explains *why* k is so high in this basin; the stability wells model shows *how high* it is and what escape requires.

Levitsky & Way (2010), *Competitive Authoritarianism*. The hybrid trap—the shallow third basin at L≈47 with 67% ten-year retention—is the feature of the stability wells model that most directly validates and extends Levitsky and Way's scholarship. Before their work, the dominant paradigm (reflected in both Polity and the earlier "transition paradigm" critiqued by Carothers, 2002) treated hybrid regimes as inherently transitional—way stations on a path toward either democracy or autocracy. Levitsky and Way demonstrated that competitive authoritarian regimes could persist for decades, featuring real but unfair elections, constrained but not eliminated opposition, and rule by law rather than rule of law. The stability wells model confirms this empirically: the hybrid trap is a *genuine attractor basin*, not a transitional slope. The GMM identifies a distinct cluster at μ=25.0 (and the broader hybrid zone at L≈20–70), and the retention rate of 67% over ten years shows that countries entering this zone tend to stay. However, the model also reveals what Levitsky and Way's qualitative framework left ambiguous: the hybrid trap is the *shallowest* of the three basins (k≈0.05 vs. k≈0.10 for democracy and k≈0.15 for tyranny). This means that while it is a genuine attractor, it is less stable than either endpoint—countries in the hybrid trap are more susceptible to shocks, more likely to transition, and more volatile in their trajectories. The implication is that Levitsky and Way were right that competitive authoritarianism is a stable regime type, but the stability wells model adds a crucial qualification: it is the *least* stable of the three equilibria, and its long-term survival depends on the absence of large shocks more than on its own structural resilience.

~~Source:~~ Analysis of 91 countries (1972-2025) · Retention = % remaining in same zone after 10 years · Basin depths calibrated to historical transition matrices from Freedom House data · GMM K=3 validates three attractor basins (μ=7.3, μ=25.0, μ=69.2)

Related Articles

[Framework Basin Overview] [Framework Escape Velocity Maps] [Attractor Basins Graphics] [Threshold Dynamics]

Figure 2. Empirical potential landscape V(L) = −log p(L). Local minima correspond to attractor basins; local maxima (saddle points) correspond to energy barriers between basins. The tyranny well at L ≈ 10 is the deepest basin; the hybrid trap at L ≈ 48 is the shallowest; the democratic plateau at L ≈ 88 is of intermediate depth. Country positions are illustrative, based on most recent observations.

6.3 Parametric Potential Models

To quantify the basin structure, we fit three parametric potential models to the empirical potential function: a single-Gaussian (parabolic) well, a double-Gaussian well, and a triple-Gaussian well. The triple-Gaussian model takes the form:

$$V(L) = -\log[w_1\phi(L; \mu_1, \sigma_1) + w_2\phi(L; \mu_2, \sigma_2) + w_3\phi(L; \mu_3, \sigma_3)] + C \tag{6}$$

Parameters are estimated by minimising the residual sum of squares (RSS) between the empirical and model potential functions, using the Nelder-Mead simplex algorithm with 30 random restarts to avoid local minima. Model selection proceeds via BIC computed from RSS.

Table 4. Parametric Potential Model Comparison

Model	Parameters	RSS	BIC	ΔBIC
Single-well (parabolic)	3	42.87	−472.3	+285.6
Double-well	7	8.14	−689.4	+68.5
Triple-well	**10**	**1.92**	**−757.9**	**0.0**

Notes: RSS computed over 200 grid points spanning L ∈ [0.5, 100]. BIC approximated from RSS: n·log(RSS/n) + k·log(n). Nelder-Mead optimisation with 30 random restarts.

The triple-well model is decisively preferred, with RSS reduced by 76% relative to the double-well model and 96% relative to the single-well model. The BIC differences are substantial (ΔBIC = 68.5 for double vs. triple, 285.6 for single vs. triple), confirming that the additional parameters of the triple-well specification are justified by the data.

6.4 Basin Properties

The estimated triple-well model reveals the following basin structure. The tyranny well, centred at L ≈ 10, is the deepest basin in the landscape. Countries that fall into this well face the highest energy barrier to escape, consistent with the empirical observation that consolidated autocracies rarely democratise without extraordinary exogenous shocks (Geddes, Wright, and Frantz 2018). The democratic plateau, centred at L ≈ 88, is an elevated basin of intermediate depth, stabilised by the redundancy of democratic institutions: free press, independent judiciary, competitive elections, and civil society organisations each independently resist erosion. The hybrid trap, centred at L ≈ 48, is the shallowest basin, consistent with the observation that countries in this zone exhibit high volatility and frequent partial transitions.

The two saddle points—the local maxima separating adjacent basins—are located at approximately L ≈ 28 (between tyranny and hybrid) and L ≈ 68 (between hybrid and democracy). These saddle points represent the energy barriers that must be overcome for a country to transition between basins. The barrier between tyranny and hybrid is higher than the barrier between hybrid and democracy, consistent with the observation that escaping consolidated autocracy is more difficult than escaping the hybrid trap.

Proposition 2 (Asymmetric Basin Depths). The three attractor basins are of unequal depth. The tyranny well is deepest, the democratic plateau is of intermediate depth, and the hybrid trap is shallowest. This asymmetry implies that the long-run stationary distribution of political regimes favours autocracy in the absence of sustained institutional investment in democratic governance.

6.5 Mean First-Passage Times

The Langevin framework permits the calculation of mean first-passage times (MFPTs)—the expected time for a country to transition from one basin to another. For an overdamped particle in a one-dimensional potential, the Kramers formula gives the MFPT from a potential minimum at position a to escape over a barrier at position b as:

$$\tau_{escape} \approx (2\pi / \omega_a \omega_b)\ exp(2\Delta V / \sigma^2) \tag{7}$$

where $\Delta V = V(b) - V(a)$ is the barrier height, and ω_a, ω_b are the curvatures of the potential at the minimum and saddle point, respectively. The exponential dependence on barrier height means that small increases in ΔV translate to large increases in escape time, explaining the remarkable persistence of consolidated autocracies and consolidated democracies alike.

7. The Critical Instability Zone

7.1 Identifying the Threshold

A central finding of our analysis is the identification of a Critical Instability Zone at L ≈ 52–55, below which the probability of democratic recovery collapses. We arrive at this estimate through three independent methods that converge on the same narrow range.

Method 1: Survival analysis. We compute the proportion of country-year observations at each liberty score that subsequently recover to L ≥ 70 within a 15-year window. This recovery rate exhibits a sharp discontinuity at L ≈ 52–55: above this threshold, approximately 82% of observations recover; below it, only 3.0% do.

Method 2: Markov transition probabilities. We estimate a discrete Markov transition matrix with 10-point liberty score bins. The transition probabilities reveal an asymmetry at L ≈ 50–55: above this level, the probability of upward transition exceeds the probability of downward transition; below it, the relationship reverses, and countries are more likely to continue declining than to recover.

Method 3: Potential landscape analysis. The empirical potential function V(L) exhibits an inflection point at approximately L ≈ 53, where the potential gradient shifts from pulling countries upward (toward the democratic plateau) to pulling them downward (toward the hybrid trap or tyranny well). This inflection point corresponds to the ridgeline between the hybrid trap basin and the democratic plateau basin in the potential landscape.

> **Result 2.** Three independent methods—survival analysis, Markov transition probabilities, and potential landscape estimation—converge on L ≈ 52–55 as a Critical Instability Zone. Below this threshold, the probability of recovery to L ≥ 70 is 3.0% (bootstrap 95% CI: 0.7–6.0%).

7.2 Bootstrap Confidence Intervals

To quantify the uncertainty around the critical threshold and the recovery probability, we conduct a nonparametric bootstrap with 1,000 iterations. In each iteration, we resample the 1,656 country-year observations with replacement, recompute the recovery rates at each liberty score level, and identify the threshold at which the recovery rate falls below 10%. The distribution of bootstrapped thresholds has a median of 53.2 and a 95% confidence interval of [50.8, 56.1]. The recovery rate below the threshold has a bootstrap distribution with median 3.0% and 95% CI [0.7%, 6.0%].

Table 5. Critical Instability Zone: Bootstrap Estimates

Parameter	Point Estimate	Bootstrap 95% CI
Threshold location	L ≈ 53	[50.8, 56.1]
Recovery rate above threshold (to L ≥ 70, 15yr)	82%	[76%, 88%]
Recovery rate below threshold	3.0%	[0.7%, 6.0%]
Odds ratio (above/below)	27.3	[12.7, 125.7]

Notes: 1,000 bootstrap resamples with replacement. Recovery defined as reaching L ≥ 70 within 15 years. CI computed as 2.5th and 97.5th percentiles of bootstrap distribution.

7.3 AR(1) Persistence

A striking feature of the data is the dominance of simple autoregressive persistence over more complex transition models. We estimate the following AR(1) specification:

$$L_{i,t+1} = \alpha + \beta L_{i,t} + \varepsilon_{i,t} \tag{8}$$

The estimated persistence parameter is $\beta = 0.96$ (SE = 0.008, $p < 0.001$), with $R^2 = 0.872$. This means that a country's liberty score in any given period is overwhelmingly determined by its score in the previous period, with only marginal contributions from other factors. The implied long-run equilibrium is $L^* = \alpha / (1 - \beta) \approx 81.6$, which falls within the democratic plateau basin identified by the GMM.

We compare the AR(1) specification against a range of stage-based transition models that assign countries to discrete categories and estimate transition probabilities between them. The AR(1) model dominates all stage-based specifications:

Table 6. Model Comparison: AR(1) vs. Stage-Based Models

Model	Parameters	R^2	AIC	ΔAIC
AR(1)	**3**	**0.872**	**8,241**	**0**
3-stage transition	9	0.714	8,892	+651
5-stage transition	25	0.742	8,784	+543
8-stage transition	56	0.769	8,703	+462

Notes: ΔAIC computed relative to the AR(1) model. Stage-based models assign countries to discrete categories and estimate Markov transition matrices between stages. All ΔAIC values exceed 300, indicating decisive preference for the AR(1) specification.

Proposition 3 (Persistence Dominance). AR(1) persistence ($\beta = 0.96$, $R^2 = 0.872$) dominates all stage-based transition models by $\Delta AIC > 300$. The underlying data-generating process for regime dynamics is better characterised by continuous drift with stochastic shocks than by discrete jumps between categorical states.

This finding does not invalidate the attractor basin framework—the GMM and potential landscape results remain robust—but it does suggest that the basins exert their influence through a continuous, gradualist mechanism rather than through discrete threshold effects. Countries do not "jump" from one basin to another; they drift slowly under the combined influence of institutional momentum ($\beta = 0.96$), mean-reverting forces (the potential gradient), and stochastic shocks. The critical instability threshold at $L \approx 52–55$ marks the point where the direction of the mean-reverting force changes sign, not a discrete barrier that countries cross.

8. Robustness and Limitations

8.1 Tyranny as Residual

The most significant structural limitation of the analysis is that Tyranny (T) is computed as a residual rather than measured independently. Since T = 100 − L − C by construction, any measurement error in L or C is mechanically transmitted to T with the opposite sign. If Freedom House systematically overstates a country's liberty score, the framework will correspondingly understate its tyranny. More subtly, the ternary constraint imposes a zero-sum structure that may not accurately reflect all political dynamics: it is at least conceivable that a country could simultaneously increase both liberty and state capacity (reducing chaos) without a corresponding increase in tyranny.

We emphasise that the ternary constraint holds *definitionally* in our framework, not as an independent empirical finding. Future research should develop independent measures of tyranny—such as indices of executive concentration, political prisoner counts, surveillance intensity, or extrajudicial violence—and test whether the ternary constraint holds empirically when all three components are measured independently.

8.2 Crosswalk Accuracy

The 67% crosswalk match rate between our derived liberty scores and the published Freedom House classification means that one-third of country-year observations exhibit non-trivial divergence from the field's standard reference point. Some of this divergence is by design: the Political Topology Index incorporates institutional erosion signals that FH updates more slowly. However, a 33% disagreement rate is substantial and warrants ongoing calibration effort.

8.3 Implementation Constraints

All statistical analysis in this paper is performed using Python's standard library (csv, math, statistics, random modules). No third-party packages (scipy, numpy, statsmodels) are used. This constraint, adopted for auditing and reproducibility purposes, limits the sophistication of available statistical methods. In particular, our bootstrap confidence intervals are computed from percentiles of the bootstrap distribution rather than from bias-corrected and accelerated (BCa) methods, and our regression standard errors assume homoskedasticity rather than using heteroskedasticity-robust (HC) estimators. We view this as an acceptable trade-off: every line of analytical code is inspectable and self-contained, which enhances reproducibility at the cost of some statistical refinement.

8.4 Temporal Coverage and Interpolation

Observations are concentrated at historical inflection points rather than distributed at regular intervals, with pre-1972 data relying heavily on V-Dem and Polity crosswalks. Linear interpolation between observation points imposes a smoothness assumption that may mask rapid within-period transitions. The time spacing varies from 1 year (for modern observations) to 10–20 years (for 19th-century observations), creating heterogeneous precision across the sample. This irregular spacing complicates the interpretation of the AR(1) parameter: $\beta = 0.96$ per period is meaningfully different depending on whether the period is one year or ten years.

8.5 Potential Endogeneity

Our analysis treats liberty scores as the dependent variable without modelling the structural determinants of regime type (economic development, inequality, ethnic fractionalization, geopolitical environment). This omission raises concerns about endogeneity: variables that simultaneously affect both the current liberty score and the probability of future transition (such as GDP per capita) may bias our estimates of persistence and transition probabilities. Acemoglu et al. (2008) demonstrated that the observed correlation between income and democracy is largely spurious once country fixed effects are included, suggesting that cross-sectional comparisons may overstate the causal role of economic development. Our panel structure provides some protection against this concern, but the irregular spacing and limited within-country variation limit our ability to include country fixed effects reliably.

8.6 Small N for Country-Specific Claims

While the pooled sample of 1,656 observations provides adequate statistical power for the GMM and potential landscape analyses, country-specific claims rest on much smaller samples. The United States, for example, has 34 observations over 225 years. Statistical inferences about individual country trajectories should be treated as illustrative rather than definitive, with confidence intervals substantially wider than those reported for the pooled sample.

8.7 Sensitivity to Classification Boundaries

The assignment of countries to basins depends on the chosen boundary values ($L < 20$ for tyranny, $L = 20–70$ for hybrid, $L > 80$ for democracy). These boundaries are derived from the GMM component means and standard deviations but inevitably involve some arbitrariness at the margins. Zone velocity estimates are sensitive to whether countries are classified by their starting-zone or ending-zone assignment: starting-zone assignment yields materially different results for the tyranny basin (e.g., +0.72/yr starting vs. −0.64/yr ending). This sensitivity should temper the interpretation of basin-specific dynamics.

9. Discussion

9.1 Relationship to Existing Theory

Our results speak directly to several unresolved debates in comparative politics.

The Levitsky-Way thesis vindicated. The hybrid trap basin provides the first formal statistical vindication of Levitsky and Way's (2010) central argument that competitive authoritarian regimes are not merely "incomplete transitions" toward democracy but constitute a distinct regime type with self-reinforcing logic. Where Levitsky and Way argued qualitatively that competitive authoritarian regimes can persist for decades through manipulated elections and constrained opposition, we quantify this persistence: the hybrid basin at $\mu \approx 47$ captures approximately 34.5% of all country-year observations and exhibits substantial temporal persistence. The GMM formally identifies it as a distinct component of the distribution, not as a noisy overlap between two components.

Acemoglu-Robinson extended to three equilibria. Our tristable framework extends the Acemoglu-Robinson (2006) model by introducing a third equilibrium that their binary choice set cannot accommodate. Their game-theoretic model predicts either full concession (democracy) or repression (dictatorship); the empirical distribution reveals a large cluster of countries at intermediate liberty scores that are neither. The hybrid trap may correspond to a "partial concession" equilibrium in which elites permit limited political opening without ceding full control—a possibility that Acemoglu, Robinson, and Verdier (2004) explored in their analysis of kleptocracy but did not integrate into the main transition framework.

Huntington's waves reinterpreted. The tristable framework offers a micro-foundation for Huntington's (1991) macro-historical pattern of democratic waves and reverse waves. Waves correspond to periods when correlated exogenous shocks (e.g., post-war settlement, decolonisation, Soviet collapse) push multiple countries over saddle points simultaneously, while reverse waves represent coordinated slides back into attractor basins. Crucially, however, the tristable model does not predict a secular trend toward democracy. The asymmetric basin structure—with the tyranny well being the deepest attractor—implies that the long-run stationary distribution may favour autocracy absent sustained institutional investment, a finding that contradicts the Whig interpretation implicit in much of the democratisation literature.

Treisman's stochastic transitions formalised. Treisman's (2020) argument that many democratisations were "by mistake"—the unintended consequences of miscalculation by authoritarian incumbents—finds a natural formalisation in the Langevin framework. Democratisation-by-mistake corresponds to a large positive stochastic shock ε pushing a country over the saddle point from the hybrid trap into the democratic basin. The model predicts that such events should be rare, stochastic, and difficult to sustain, exactly the pattern Treisman documents empirically.

9.2 The Democratic Plateau as Engineered Stability

A notable feature of the potential landscape is that the democratic plateau is *not* the deepest basin. In the absence of institutional reinforcement, the natural resting state of the system—the global minimum of the potential function—would be the tyranny well. The stability of consolidated democracies is therefore not a natural equilibrium but an *engineered* one, maintained by the redundancy of institutional checks: when one institution is weakened, others compensate. Free press monitors judicial capture; courts check executive overreach; competitive elections allow peaceful replacement of incumbents; civil society organisations mobilise opposition to norm violations. The institutional ecology of democracy creates a self-reinforcing dynamic—but only so long as a sufficient number of these institutions remain functional.

This interpretation has a direct implication for the critical instability threshold. The threshold at L ≈ 52–55 may correspond to the point at which enough institutional checks have been degraded that the remaining institutions can no longer compensate. Below this point, the erosion of any one institution accelerates the erosion of all others, creating a positive feedback loop that drives the country downward through the hybrid trap toward the tyranny well. This is consistent with the eight-step erosion model (norm erosion, information capture, judicial capture, legislative subordination, regulatory capture, civil society suppression, electoral manipulation, constitutional consolidation), in which judicial capture (Step 3) is identified as the critical inflection point because an independent judiciary is the institution most central to the maintenance of all other institutional constraints.

9.3 Implications for Early-Warning Systems

The identification of a quantitative critical instability threshold has practical implications for the design of democratic early-warning systems. Current early-warning approaches (e.g., Goldstone et al. 2010; Hegre et al. 2013) rely on structural risk factors (GDP, ethnic fractionalization, neighbourhood effects) that change slowly and are difficult to modify through intervention. Our framework suggests that monitoring the *liberty score trajectory*—specifically, whether a country is approaching the L ≈ 52–55 threshold from above—may provide a more actionable early-warning signal. The finding that recovery probability collapses from approximately 82% to 3% as countries cross this threshold implies that interventions to arrest democratic erosion are dramatically more effective when undertaken early, before the threshold is crossed.

9.4 The Role of Economic Development

Our analysis does not directly model the relationship between economic development and regime type, but the literature suggests important interactions. Przeworski et al. (2000) demonstrated that no democracy above approximately $6,000 per capita GDP (in 1985 PPP dollars) has ever collapsed. More recently, Treisman (2015) has argued that the income-democracy relationship is more nuanced than previously supposed, with economic modernisation creating both pro-democratic and pro-authoritarian pressures. The implication for our framework is that the depth and location of the attractor basins may vary with economic development: wealthier countries may have deeper democratic plateaus and shallower tyranny wells, effectively raising the energy barriers against democratic collapse. This is a promising direction for future research that would require incorporating GDP per capita as a time-varying covariate in the potential landscape estimation.

10. Conclusion

This paper has presented evidence that the political regime landscape is governed by tristable dynamics with three distinct attractor basins: a democratic plateau ($L > 80$), a hybrid trap ($L \approx 20$–70), and a tyranny well ($L < 20$). Using a 225-year panel of 91 countries ($N = 1{,}656$), we have shown that:

1. A three-component Gaussian Mixture Model is decisively preferred by BIC over two-component and single-component alternatives ($\Delta BIC = 193.9$ over $K = 2$), providing formal statistical evidence for three rather than two attractor basins.

2. The empirical potential function $V(L) = -\log p(L)$, derived from a Langevin SDE framework, exhibits three local minima of unequal depth. The tyranny well is the deepest attractor, suggesting that autocracy is the natural resting state absent sustained institutional investment in democratic governance.

3. A Critical Instability Zone exists at $L \approx 52$–55, below which the probability of democratic recovery collapses to 3.0% (95% CI: 0.7–6.0%). Three independent estimation methods—survival analysis, Markov transition probabilities, and potential landscape inflection—converge on this range.

4. AR(1) persistence ($\beta = 0.96$) dominates all stage-based transition models ($\Delta AIC > 300$), implying that regime dynamics are governed by continuous drift rather than discrete threshold effects, though the attractor basin structure shapes the drift direction.

These findings challenge the dominant bistable framework in comparative politics and suggest that the hybrid regime zone is not a transient way-station between democracy and autocracy but a distinct attractor with its own stabilising dynamics. The practical implication is sobering: countries that slip below the critical instability threshold face overwhelming odds against recovery, and interventions to arrest democratic erosion are far more effective when undertaken before the threshold is crossed than after.

We acknowledge significant limitations: Tyranny is measured as a residual rather than independently, the crosswalk between data sources achieves only 67% accuracy, the analysis is implemented using only Python standard library functions, and country-specific inferences rest on small samples. These limitations constrain the precision of our quantitative estimates without, we argue, undermining the qualitative finding that the regime landscape is tristable rather than bistable.

Future research should pursue four priorities. First, independent measurement of Tyranny would permit an empirical test of the ternary constraint rather than imposing it by construction. Second, a continuous-time Markov chain (CTMC) estimated via matrix exponentials would provide a unified framework that bridges the discrete transition data and the continuous Langevin dynamics. Third, the incorporation of time-varying covariates (GDP per capita, inequality, democratic tenure) would allow the basin structure to vary across countries and periods. Fourth, external validation against the forthcoming 2026 reports from Freedom House and V-Dem would test whether the critical instability threshold identified here has predictive power for ongoing cases of democratic erosion.

References

Acemoglu, D., Johnson, S., Robinson, J. A., and Yared, P. (2008). Income and Democracy. *American Economic Review*, 98(3), 808–842.

Acemoglu, D. and Robinson, J. A. (2006). *Economic Origins of Dictatorship and Democracy*. Cambridge: Cambridge University Press.

Acemoglu, D., Robinson, J. A., and Verdier, T. (2004). Kleptocracy and Divide-and-Rule: A Model of Personal Rule. *Journal of the European Economic Association*, 2(2–3), 162–192.

Bermeo, N. (2016). On Democratic Backsliding. *Journal of Democracy*, 27(1), 5–19.

Boix, C., Miller, M., and Rosato, S. (2013). A Complete Data Set of Political Regimes, 1800–2007. *Comparative Political Studies*, 46(12), 1523–1554.

Castellano, C., Fortunato, S., and Loreto, V. (2009). Statistical Physics of Social Dynamics. *Reviews of Modern Physics*, 81(2), 591–646.

Coppedge, M., Gerring, J., Knutsen, C. H., et al. (2024). V-Dem Codebook v14. *Varieties of Democracy (V-Dem) Project*.

Dahl, R. A. (1971). *Polyarchy: Participation and Opposition*. New Haven: Yale University Press.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society: Series B*, 39(1), 1–38.

Diamond, L. (2002). Thinking about Hybrid Regimes. *Journal of Democracy*, 13(2), 21–35.

Diamond, L. (2015). Facing Up to the Democratic Recession. *Journal of Democracy*, 26(1), 141–155.

Freedom House. (2025). *Freedom in the World 2025*. Washington, DC: Freedom House.

Fund for Peace. (2024). *Fragile States Index 2024*. Washington, DC: Fund for Peace.

Galam, S. (2012). *Sociophysics: A Physicist's Modeling of Psycho-political Phenomena*. New York: Springer.

Geddes, B., Wright, J., and Frantz, E. (2018). *How Dictatorships Work*. Cambridge: Cambridge University Press.

Goldstone, J. A., Bates, R. H., Epstein, D. L., et al. (2010). A Global Model for Forecasting Political Instability. *American Journal of Political Science*, 54(1), 190–208.

Haggard, S. and Kaufman, R. R. (2021). *Backsliding: Democratic Regress in the Contemporary World*. Cambridge: Cambridge University Press.

Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57(2), 357–384.

Hegre, H., Karlsen, J., Nygard, H. M., Strand, H., and Urdal, H. (2013). Predicting Armed Conflict, 2010–2050. *International Studies Quarterly*, 57(2), 250–270.

Huntington, S. P. (1991). *The Third Wave: Democratization in the Late Twentieth Century*. Norman: University of Oklahoma Press.

Kramers, H. A. (1940). Brownian Motion in a Field of Force and the Diffusion Model of Chemical Reactions. *Physica*, 7(4), 284–304.

Levitsky, S. and Way, L. A. (2002). The Rise of Competitive Authoritarianism. *Journal of Democracy*, 13(2), 51–65.

Levitsky, S. and Way, L. A. (2010). *Competitive Authoritarianism: Hybrid Regimes After the Cold War*. Cambridge: Cambridge University Press.

Levitsky, S. and Ziblatt, D. (2018). *How Democracies Die*. New York: Crown.

Marshall, M. G. and Gurr, T. R. (2020). Polity5: Political Regime Characteristics and Transitions, 1800–2018. *Center for Systemic Peace*.

McLachlan, G. J. and Peel, D. (2000). *Finite Mixture Models*. New York: Wiley.

Przeworski, A. (2000). *Democracy and Development: Political Institutions and Well-Being in the World, 1950–1990*. Cambridge: Cambridge University Press.

Przeworski, A. and Limongi, F. (1997). Modernization: Theories and Facts. *World Politics*, 49(2), 155–183.

Schedler, A. (2006). *Electoral Authoritarianism: The Dynamics of Unfree Competition*. Boulder: Lynne Rienner.

Scheffer, M., Bascompte, J., Brock, W. A., et al. (2009). Early-Warning Signals for Critical Transitions. *Nature*, 461(7260), 53–59.

Schumpeter, J. A. (1942). *Capitalism, Socialism, and Democracy*. New York: Harper.

Schwarz, G. (1978). Estimating the Dimension of a Model. *Annals of Statistics*, 6(2), 461–464.

Sornette, D. (2003). *Why Stock Markets Crash: Critical Events in Complex Financial Systems*. Princeton: Princeton University Press.

Svolik, M. W. (2012). *The Politics of Authoritarian Rule*. Cambridge: Cambridge University Press.

Svolik, M. W. (2019). Polarization versus Democracy. *Journal of Democracy*, 30(3), 20–32.

Treisman, D. (2015). Income, Democracy, and Leader Turnover. *American Journal of Political Science*, 59(4), 927–942.

Treisman, D. (2020). Democracy by Mistake: How the Errors of Autocrats Trigger Transitions to Freer Government. *American Political Science Review*, 114(3), 792–810.

van Kampen, N. G. (2007). *Stochastic Processes in Physics and Chemistry*, 3rd ed. Amsterdam: North-Holland.

V-Dem Institute. (2025). V-Dem Annual Democracy Report 2025. *University of Gothenburg*.

World Bank. (2023). *Worldwide Governance Indicators 2023*. Washington, DC: World Bank.

Appendix A. Technical Details

A.1 EM Algorithm Implementation

The Expectation-Maximisation algorithm for GMM fitting proceeds as follows. Given K components and N observations $\{L_1, ..., L_N\}$:

Initialisation. Component means are initialised using K-means++ style seeding: the first mean is drawn uniformly at random from the data; subsequent means are drawn with probability proportional to the squared distance from the nearest existing mean. Component standard deviations are initialised to $2\sigma_X/K$, where σ_X is the overall standard deviation. Mixing weights are initialised to 1/K.

E-step. For each observation i and component k, compute the responsibility:

$$\gamma_{ik} = \pi_k\, \phi(L_i; \mu_k, \sigma_k) \,/\, \Sigma_{j=1}^{K}\, \pi_j\, \phi(L_i; \mu_j, \sigma_j) \tag{A1}$$

M-step. Update parameters using the responsibilities:

$$N_k = \Sigma_i\, \gamma_{ik}, \quad \pi_k = N_k/N, \quad \mu_k = \Sigma_i\, \gamma_{ik}\, L_i \,/\, N_k, \quad \sigma_k^{2} = \Sigma_i\, \gamma_{ik}(L_i - \mu_k)^{2} \,/\, N_k \tag{A2}$$

Convergence. Iterate until the change in log-likelihood between successive iterations falls below $\varepsilon = 10^{-6}$, or until a maximum of 300 iterations is reached.

Dead component handling. If the effective sample size N_k for any component falls below 10^{-10}, the component is reinitialised by drawing a new mean uniformly from the data and resetting its standard deviation to the overall standard deviation.

Multiple restarts. For K > 1, the algorithm is run 20 times with different random seeds, and the solution with the highest log-likelihood is retained.

A.2 Bootstrap Procedure

Bootstrap confidence intervals are constructed as follows:

1. Draw a bootstrap sample of size N with replacement from the original N = 1,656 observations.

2. Re-estimate the K = 3 GMM on the bootstrap sample, initialising near the full-sample solution (adding Gaussian noise with $\sigma = 2$ to each mean) to improve convergence.

3. Sort the estimated components by mean to ensure consistent labelling across bootstrap iterations.

4. Record the estimated parameters (π_k, μ_k, σ_k) for each component.

5. Repeat steps 1–4 for B = 1,000 iterations.

6. Report the 2.5th and 97.5th percentiles of the bootstrap distribution as the 95% confidence interval.

The same bootstrap procedure is applied to the potential function estimation (re-estimating KDE and fitting the triple-well model on each bootstrap sample) and to the critical instability threshold estimation (recomputing recovery rates on each bootstrap sample).

A.3 Potential Function Estimation

The nonparametric potential function is estimated as follows:

1. **KDE.** A Gaussian kernel density estimate is computed at 200 evenly spaced grid points on [0.5, 100] using bandwidth $h = \max(h_{Silverman}, 3.0)$, where $h_{Silverman} = 0.9 \cdot \min(\sigma_X, IQR/1.34) \cdot N^{-0.2}$.

2. **Potential.** The empirical potential is computed as $V(L) = -\log p(L)$, with a floor of $p = 10^{-10}$ to prevent numerical overflow. The potential is shifted so that $\min(V) = 0$.

3. **Parametric fitting.** The triple-Gaussian potential model (Equation 6) is fitted by minimising RSS using the Nelder-Mead simplex algorithm (parameters: reflection $\alpha = 1.0$, expansion $\gamma = 2.0$, contraction $\rho = 0.5$, shrink $\sigma = 0.5$; maximum 5,000 iterations; convergence tolerance 10^{-10}) with 30 random restarts. Parameter bounds enforce positive weights and standard deviations and restrict component means to plausible ranges (tyranny: [0, 30]; hybrid: [25, 65]; democracy: [60, 100]).

A.4 Replication

All analysis scripts, data files, and replication instructions are publicly available. The GMM estimation is implemented in `b1_gmm_model_comparison.py` and the potential function estimation in `b2_potential_function.py`. Both scripts require only the Python standard library (version 3.8 or later) and the input data file `political-topology-flat.csv`. The complete dataset of 91 countries, 225 years, and 1,656 observations is available in both CSV and XLSX formats.

Author Note. Cambridge Governance Labs is an independent research organisation based in Cambridge, United Kingdom, focused on quantitative analysis of governance systems. The Political Topology project is an ongoing

research programme that develops open-source tools and datasets for the study of political regime dynamics. All data, code, and methodological documentation are freely available. The authors declare no conflicts of interest. We thank three anonymous referees for comments on an earlier draft and acknowledge the foundational data contributions of Freedom House, the V-Dem Institute, the Fund for Peace, and the World Bank.

Data Availability. The complete dataset and all replication scripts are available at the Cambridge Governance Labs Political Topology repository. No proprietary data or restricted-access materials are required for replication.