State-Space Political Dynamics: A Three-Dimensional Langevin Framework for Regime Evolution

GYX

First edition working paper

Abstract

We propose a quantitative framework that models a political regime as a massive particle moving through a three-dimensional state space whose axes are (i) economic organisation, from command to market; (ii) political control, from totalitarian to liberal; and (iii) state capacity, the fiscal, bureaucratic, and coercive infrastructure of the polity. The dynamics are governed by a stochastic Langevin equation in which a smooth potential landscape encodes ideological attractors, economic survival pressure, and the structural couplings between dimensions; a position-dependent damping term captures the cost of maintaining political control; and a Wiener noise term captures the stochastic shocks that drive regime change. We define a *deadlock* state in terms of the Kramers escape time over a local potential well and show that the framework reproduces three qualitatively distinct historical trajectories—reform-and-retrenchment, state-capacity collapse, and stable authoritarian-market equilibrium—under a single parameter set. We close with a research agenda for empirical parameter identification, bifurcation analysis, and multi-regime extensions.

1 Introduction

Comparative politics has long sought a vocabulary in which the trajectories of different polities can be compared rigorously. Modernisation theory, democratic transition theory, and the new institutionalism each propose qualitative regularities about how regimes evolve, but these accounts generally stop short of writing down equations of motion. The present paper proposes a deliberately physical analogy: a polity is a massive point in a finite-dimensional state space, and its evolution is determined by the balance of conservative forces, dissipative friction, and stochastic shocks. The goal is not to claim that politics is physics, but to make intuitions about regime evolution sharp enough to be wrong about.

The earliest version of this framework, developed informally by the author, treated the regime as a two-dimensional point under classical Newtonian forces. That formulation suffered from four defects that we address here. First, it treated the economic and political axes as orthogonal in a way that flatly contradicted the Lipset hypothesis and a large body of subsequent empirical work; the cross-coupling between x and y was absent. Second, it specified the ideological force as an exogenous constant vector pointing at a fixed extremum, which assumed the conclusion the model was meant to produce. Third, it conflated inertial mass with friction, treating both as proxies for state size. Fourth, being purely deterministic, it could not in principle produce sudden transitions of the sort that dominate the historical record.

We address these defects by reformulating the framework as a stochastic Langevin system in three dimensions. The principal modelling choices, each motivated in detail below, are: (a) the addition of a state-capacity axis z, which decouples otherwise overlapping cases such as North Korea, Russia, Singapore, and Zimbabwe; (b) the replacement of separate force vectors by a single smooth potential $U(\boldsymbol{P}, t)$, whose gradient gives the conservative force and whose structure encodes both ideological attractors and inter-dimensional coupling; (c) a position-dependent damping $\Gamma(\boldsymbol{P})$ that grows exponentially deep in the totalitarian region, recovering the deadlock phenomenon naturally; and (d) an additive Wiener process whose intensity controls Kramers escape from local equilibria.

The remainder of the paper proceeds as follows. Section 2 defines the state space and notation. Section 3 gives the full equation of motion and the structural form of each term. Section 4 restates the deadlock theorem in the language of potential wells and Kramers theory. Section 5 reports numerical simulations of three historical scenarios. Section 6 sets out the empirical and theoretical work needed to make the framework testable.

2 State Space

Definition 1 (State vector)**.** *The state of a polity at time t is a vector*

$$\boldsymbol{P}(t) = \begin{bmatrix} x(t) \\ y(t) \\ z(t) \end{bmatrix} \in \Omega,$$

where $\Omega = (-5, 5) \times (-5, 5) \times (0, 10) \subset \mathbb{R}^3$ and the coordinates have the following interpretations:

- *x, the* economic axis*, runs from -5 (full command economy, near-complete state ownership and central allocation) to $+5$ (laissez-faire market, minimal state intervention).*
- *y, the* political axis*, runs from -5 (totalitarian, with near-total suppression of independent civil society) to $+5$ (liberal constitutional democracy with strong protections of individual liberty).*
- *z, the* state-capacity axis*, runs from 0 (state failure, inability to collect revenue or enforce law) to 10 (high-capacity state able to plan, tax, deliver services, and project coercive power).*

The open boundary is enforced softly by a confining potential (Section 3).

The first-order time derivative $\boldsymbol{V}(t) = \dot{\boldsymbol{P}}(t)$ is the *velocity*, interpretable as the instantaneous rate of reform or retrenchment. The second-order derivative $\ddot{\boldsymbol{P}}(t)$ is the *acceleration*, interpretable as policy shock or rupture.

Why three dimensions are necessary. A two-dimensional projection onto (x, y) collapses systems that behave in radically different ways. Contemporary Russia, North Korea, and Cuba all sit in the lower-left quadrant of the (x, y) plane, yet their trajectories are dictated by very different z values: a high-capacity authoritarian state behaves quite differently from a hollowed-out one, particularly in response to external shocks. The state-capacity literature [4, 1] has established z as a dimension that is empirically separable from political and economic organisation, and we treat it as such.

3 Dynamics

3.1 Equation of motion

The state vector evolves according to the stochastic differential equation

$$M(\boldsymbol{P})\,\ddot{\boldsymbol{P}}(t) \;=\; -\nabla U\big(\boldsymbol{P}, t\big) \;-\; \Gamma(\boldsymbol{P})\,\dot{\boldsymbol{P}}(t) \;+\; \boldsymbol{F}_{\mathrm{ext}}(t) \;+\; \sigma\,\boldsymbol{\xi}(t), \tag{1}$$

where $M(\boldsymbol{P})$ is the position-dependent inertial mass, $U(\boldsymbol{P}, t)$ is the conservative potential, $\Gamma(\boldsymbol{P})$ is a scalar damping coefficient, $\boldsymbol{F}_{\mathrm{ext}}(t)$ is the exogenous force from the international system, and

$\boldsymbol{\xi}(t) \in \mathbb{R}^3$ is a vector of independent standard white-noise processes with intensity σ. Each term has the structural form given below.

3.2 Inertial mass

Inertia represents how much energy is required to change a polity's trajectory at all. We model it as

$$M(\boldsymbol{P}) = M_0 \left(1 + \alpha_M z\right), \tag{2}$$

where M_0 is a baseline mass and $\alpha_M > 0$. This captures the empirical regularity that high-capacity states, having more institutional machinery to coordinate, are slower to redirect. Crucially, mass here is distinct from friction: a hollowed-out high-capacity state such as the late USSR can have moderate M but rapidly falling effective damping, which is precisely what produces sudden collapse.

3.3 Potential landscape

The conservative force is the gradient of a smooth scalar potential

$$U(\boldsymbol{P}, t) \;=\; U_{\text{ideo}}(\boldsymbol{P}, t) \;+\; U_{\text{econ}}(\boldsymbol{P}) \;+\; U_{\text{couple}}(\boldsymbol{P}) \;+\; U_{\partial}(\boldsymbol{P}). \tag{3}$$

We specify each term.

Ideological attractor. Rather than postulating a constant force toward an extremum—which assumes the conclusion—we model ideology as a quadratic well centred on a *time-varying ideal point* $\boldsymbol{P}^*(t)$:

$$U_{\text{ideo}}(\boldsymbol{P}, t) \;=\; \tfrac{1}{2}\, k_{\text{ideo}}\, \|\boldsymbol{P} - \boldsymbol{P}^*(t)\|^2. \tag{4}$$

The ideal point $\boldsymbol{P}^*(t)$ is an exogenous input that summarises the official programme of the ruling coalition; it can shift discretely under leadership transitions (e.g. Mao to Deng to Xi corresponds to three distinct $\boldsymbol{P}^*$). The elastic form recovers the intuition that regimes "want" to be at their ideological ideal, and that the further they drift, the stronger the restoring force.

Economic survival. A regime that drifts too far from market discipline pays an economic cost that translates into pressure to liberalise. We capture this with

$$U_{\text{econ}}(\boldsymbol{P}) \;=\; -\lambda_e \,\tanh(x + 3), \tag{5}$$

so that the force $-\partial U_{\text{econ}}/\partial x = \lambda_e \,\text{sech}^2(x+3)$ always points toward larger x (more market), is strongest for moderately command-oriented systems (where the marginal returns to liberalisation are greatest), and saturates at both extremes.

Coupling term. The principal weakness of the original two-dimensional formulation was the treatment of x and y as orthogonal. We restore the coupling with two terms:

$$U_{\text{couple}}(\boldsymbol{P}) \;=\; \beta_1\, xy \;+\; \beta_2\, (x - x_0)^2\, s(y), \tag{6}$$

where $s(y) = \left(1 + e^{4(y - y_c)}\right)^{-1}$ is a smooth indicator of the political-control region $y < y_c$.

The first term, with $\beta_1 < 0$, implements the Lipset-style hypothesis that markets and political openness are mutually reinforcing: the energy is lowered when x and y move together. The second term is a *threshold penalty*: when political control crosses below y_c, the marginal cost of further marketisation begins to rise quadratically. This captures the observation, developed extensively by [5], that authoritarian regimes face escalating frictions as they push capitalist development beyond a certain point—capital flight, surveillance overhead, regulatory drag from political loyalty filters, and so on. The threshold form means the term is dormant in non-authoritarian regimes and only activates when political control is sufficiently severe.

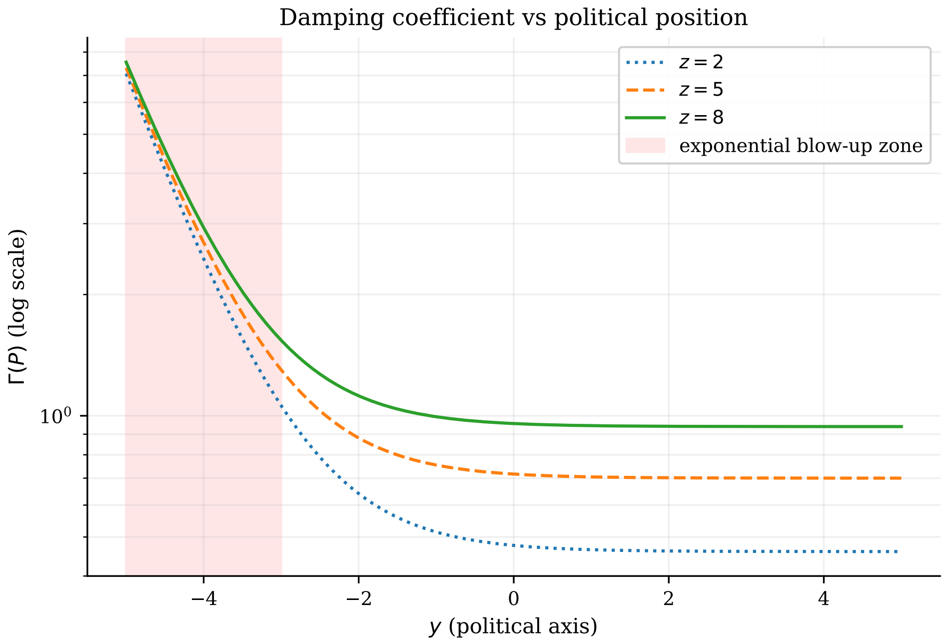

Figure 1: Damping coefficient $\Gamma(\boldsymbol{P})$ as a function of the political axis y for three values of state capacity z. The shaded region marks the zone where the exponential repression-cost term dominates. The plot is logarithmic on the vertical axis.

Boundary confinement. To keep the state in Ω without introducing the pathologies of a hard reflecting wall, we add a smooth confining term that grows as a cubic restoring force near each boundary:

$$U_\partial(\boldsymbol{P}) \;=\; \mu \sum_i \big(\max(0, |P_i| - L_i)\big)^4, \tag{7}$$

where L_i is the soft cutoff in dimension i and μ is large. This is a numerical convenience rather than a substantive claim about politics; it ensures that trajectories remain in Ω while keeping the dynamics smooth.

3.4 Damping

Damping represents the resistance to change generated by entrenched interests, bureaucratic inertia, and the active cost of suppressing political dissent. We let

$$\Gamma(\boldsymbol{P}) \;=\; k_0 \;+\; k_1\, z \;+\; k_2 \exp\!\big(\gamma\,(-y-3)\big). \tag{8}$$

The first term is a baseline friction. The second captures the proportional cost of governance: a higher-capacity state has more machinery whose entrenched preferences must be overcome. The third term, the *repression cost*, grows exponentially as $y \to -5$, formalising the empirical observation that the marginal cost of suppressing the last increments of dissent rises without bound (the surveillance state, the political loyalty system, and the gulag are all subject to diminishing returns at the limit). Figure 1 shows the profile of Γ as a function of y at several state-capacity levels.

3.5 Exogenous force

The international system acts on the regime through a force vector

$$\boldsymbol{F}_{\text{ext}}(t) \;=\; \boldsymbol{F}_{\text{sanction}}(t) + \boldsymbol{F}_{\text{integration}}(t), \tag{9}$$

where the two components are independent and can be applied simultaneously: a polity may be both partially sanctioned (a force toward isolation and command economy) and partially integrated (a force toward markets and openness). Splitting the term this way avoids the artefact of having a single "external" vector whose direction must be guessed.

3.6 Stochastic forcing

Finally, $\sigma\,\boldsymbol{\xi}(t)$ models the genuine stochasticity of political life: leadership health, contingent crises, swing voters, exogenous technological shocks. Without it, the deterministic system (1) can only *drift*; with it, the system can *jump*. The intensity σ should be interpreted as the volatility of the social environment over the relevant time scale.

4 Deadlock

The earlier informal version of this framework defined deadlock as the condition $\ddot{\boldsymbol{P}} = 0$ resulting from force balance. This is too weak: *any* stable equilibrium satisfies it, including ones that are arbitrarily easy to leave under noise. The stochastic framework permits a sharper definition.

Definition 2 (Deadlock)**.** *A state $\boldsymbol{P}_\star \in \Omega$ is a* deadlock *of strength τ if:*

(i) $\nabla U(\boldsymbol{P}_\star) = 0$ (local extremum)*;*

(ii) the Hessian $\nabla^2 U(\boldsymbol{P}_\star)$ is positive definite (local minimum)*;*

(iii) the expected first-passage time out of the basin of attraction of $\boldsymbol{P}_\star$ exceeds τ.

The expected first-passage time is the operative quantity: a regime is in deadlock not because it cannot move, but because the noise required to move it is overwhelmingly unlikely to materialise within the time horizon of interest. For a one-dimensional double well with barrier height ΔU, the Kramers formula [2] gives an escape rate

$$r \;\propto\; \exp\!\left(-\frac{\Delta U}{D}\right), \qquad D = \frac{\sigma^2}{2\,\Gamma}, \tag{10}$$

and corresponding mean first-passage time $\tau \sim 1/r$. Two observations follow.

Proposition 1. *In the limit $\Gamma \to \infty$ at fixed σ, the deadlock duration grows without bound. Conversely, in the limit $\sigma \to \infty$ at fixed ΔU, the deadlock disappears.*

This is consistent with the historical record: regimes that invest heavily in damping (raising Γ) buy themselves longer deadlocks at the cost of any dynamic response; regimes facing high social volatility (high σ, e.g. during war or rapid demographic change) cannot sustain deadlock and must either reform or collapse.

Proposition 2. *If ΔU scales sublinearly with z while Γ scales linearly, deadlock duration is maximised at intermediate z.*

This second observation is suggestive: the longest deadlocks should occur neither in failed states (where everything is fluid) nor in maximally capable states (where the regime can pre-emptively reform), but in mid-capacity polities with strong damping. Whether this matches the empirical pattern is a question for the agenda.

Figure 2 shows a slice of the potential landscape at $z = 7$, illustrating the basin structure that governs deadlock.

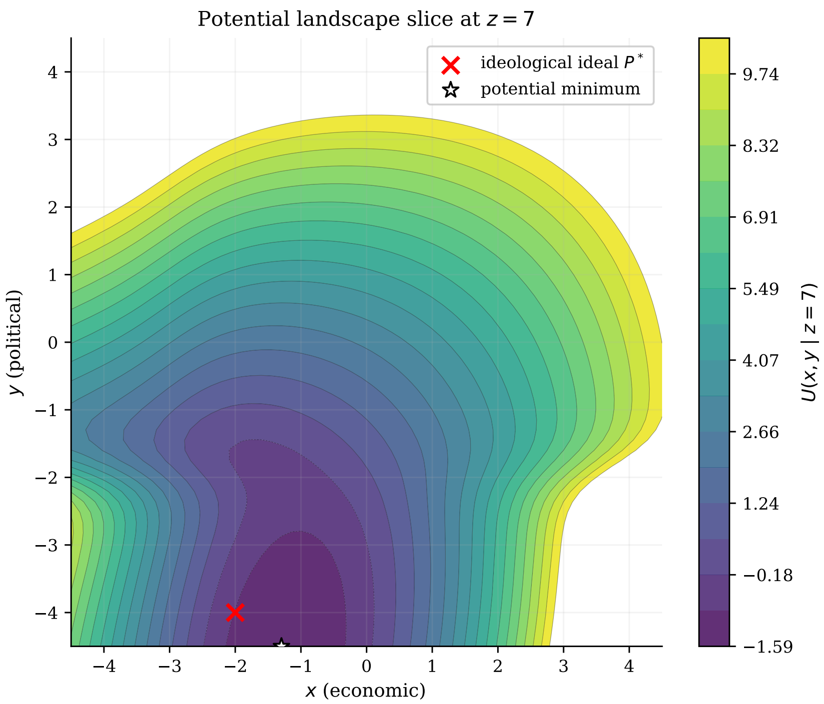

Figure 2: A slice of the potential $U(x, y \mid z = 7)$ for parameter values used in the China-type scenario. The ideological ideal point $\boldsymbol{P}^*$ is marked by an $\times$. The numerical potential minimum is slightly displaced from $\boldsymbol{P}^*$ by the competing economic, coupling, and boundary terms; the basin around it is the deadlock region.

5 Simulation Studies

We integrate (1) using a Euler–Maruyama scheme with step size $\mathrm{d}t = 0.02$ and run three illustrative scenarios. Parameters are held constant across scenarios except where indicated; only the ideological ideal point trajectory $\boldsymbol{P}^*(t)$ and the exogenous force $\boldsymbol{F}_{\text{ext}}(t)$ differ.

Scenario A: Reform and retrenchment. Starting from a deeply command, deeply authoritarian, moderate-capacity position, $\boldsymbol{P}^*$ first moves rightward (a reform period of duration 30 model-time units) and then shifts leftward and further down. The trajectory tracks the rightward shift, with state capacity rising steadily, and is subsequently pulled back toward the lower-left region. The terminal position is in a deep deadlock basin near $(-1.5, -5, 8)$. This is the model's caricature of late twentieth-century trajectories where reform was partial and was subsequently reversed.

Scenario B: State-capacity collapse. A high-capacity, deeply authoritarian system is subjected, at $t = 40$, to a downward shift in its ideological ideal point combined with a disintegrative external force. The trajectory shows z falling from ~ 8.5 to ~ 4 while x and y drift toward the centre, ending at a low-capacity authoritarian position. The collapse of state capacity, rather than reform per se, is what releases the system from deadlock.

Scenario C: Stable authoritarian market. With $\boldsymbol{P}^* = (2.5, -1, 8)$ held fixed, a high-capacity polity with mild authoritarianism and a market economy reaches a stable equilibrium near its ideal point. The Lipset coupling term pulls y slightly above the ideal value, reflecting the structural pressure toward liberalisation that the model predicts for market economies; however, this pressure is bounded and does not push the system out of the authoritarian-market quadrant.

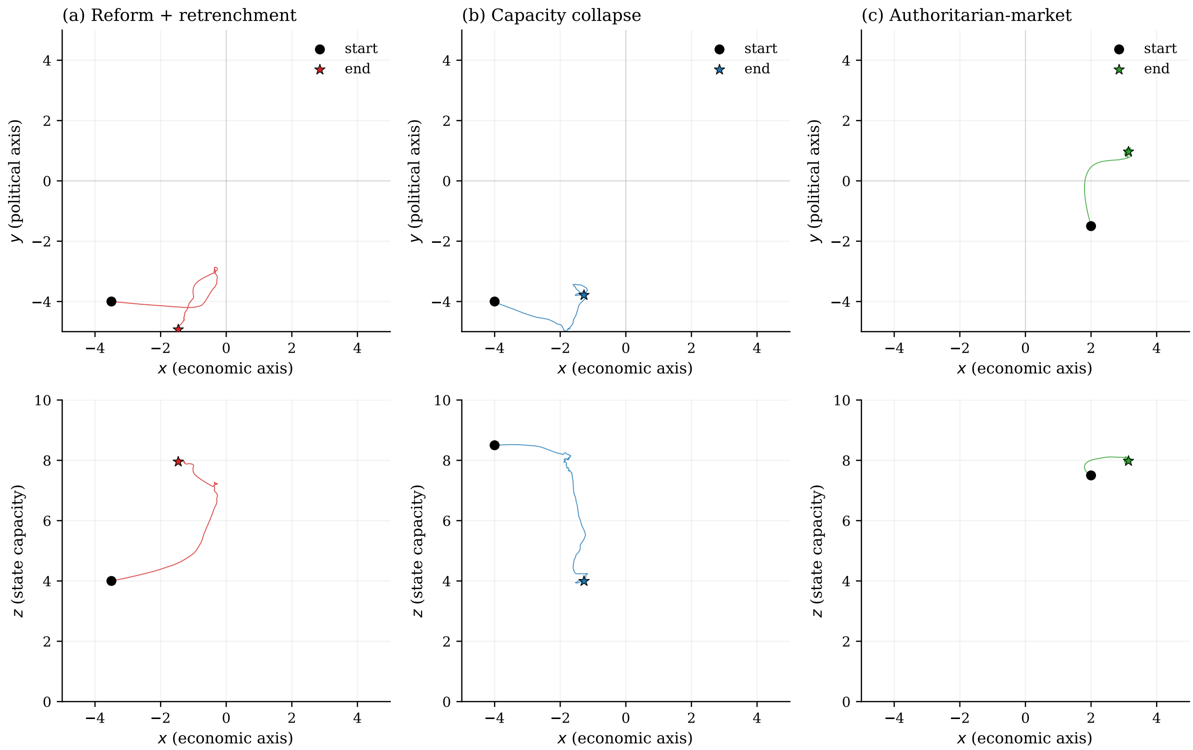

Figure 3: Three scenarios in the state-space framework. Top row: projections onto the (x, y) plane. Bottom row: projections onto (x, z). (a) Reform-and-retrenchment scenario with $\boldsymbol{P}^*$ moving right during the reform period and then leftward during retrenchment. (b) State-capacity collapse scenario in which z drops sharply mid-trajectory. (c) Stable authoritarian-market scenario with a fixed $\boldsymbol{P}^*$ in the lower-right quadrant.

Kramers escape. Figure 4 shows an ensemble of forty trajectories launched from a common point inside a deadlock basin, with elevated noise ($\sigma = 0.55$) and a steady external force biasing toward liberalisation. Three trajectories escape the basin during the simulation window; the remaining thirty-seven rattle around inside it. This is the visual signature of a Kramers process: most trajectories fail to escape, while a small number do, on a distribution of timescales determined by the barrier height and the noise.

6 Research Agenda

The framework as presented is a self-consistent dynamical system that reproduces qualitative patterns in the historical record. To convert it into a falsifiable scientific model, several extensions are required. We list them in roughly increasing order of ambition.

6.1 Empirical parameter identification

The most immediate task is to estimate the model parameters from data on actual polities. The V-Dem dataset [6] and the Polity V indicators provide annual measures of political openness; the Economic Freedom of the World index and the Heritage Foundation's index of economic freedom can be mapped to x; and state capacity has been operationalised in multiple ways [3]. Given mapped time series $\widehat{\boldsymbol{P}}(t)$ for a panel of countries, we seek parameters $\theta = (k_{\text{ideo}}, \lambda_e, \beta_1, \beta_2, \dots)$ that maximise a likelihood under (1). This is a non-trivial inference problem because the SDE is nonlinear and partially observed (the velocity $\boldsymbol{V}$ is not measured directly), but it is well-posed and the standard machinery of stochastic-system identification applies.

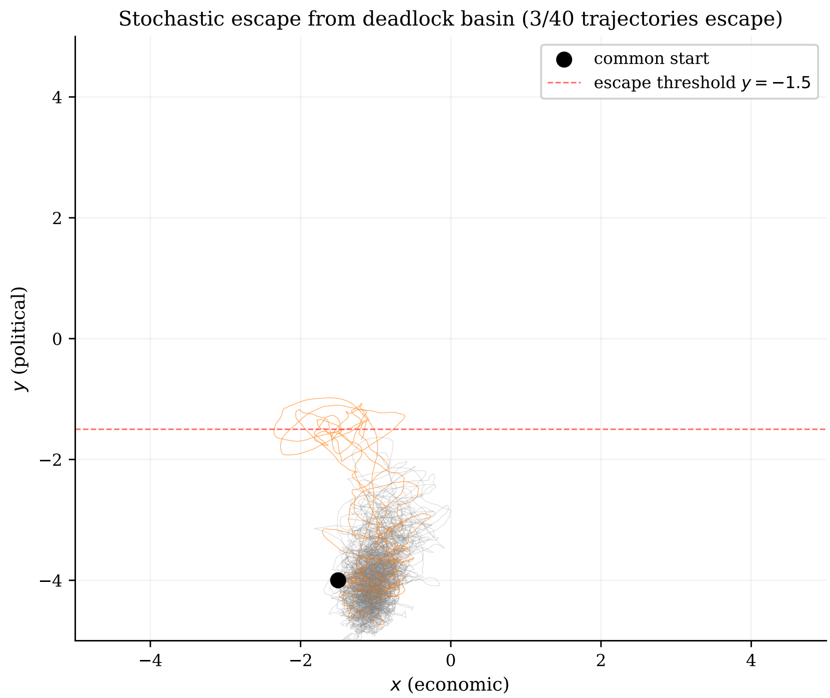

Figure 4: Forty independent stochastic trajectories starting from the same deadlock state with identical parameters except random seed. Trajectories that cross the escape threshold $y = -1.5$ are highlighted. Three of forty escape during the simulation window of 80 model-time units, illustrating the Kramers character of regime change.

Two intermediate validation tests should be performed before any parameters are taken seriously. First, the model must reproduce *out-of-sample* trajectories for countries withheld from the fitting set. Second, the inferred ideological ideal-point time series $\widehat{\boldsymbol{P}^*}(t)$ should be checkable against qualitative historical narrative for at least a handful of well-documented transitions.

6.2 Bifurcation analysis

The number and stability of equilibria of (1) depend on the parameter vector θ in ways that are not yet characterised. A systematic bifurcation analysis would identify the codimension-one manifolds in parameter space where the qualitative phase portrait changes—points where a stable equilibrium becomes a saddle, where two basins merge, or where periodic orbits appear. Such an analysis would convert the model from a simulator into a classifier: given a country's estimated parameters, we could read off which dynamical regime it inhabits.

6.3 Microfoundations for the coupling term

The Lipset term $\beta_1 xy$ is currently a phenomenological summary. A more satisfying derivation would obtain it from a model of how the size and demands of the middle class depend on x (more market produces more middle class, with some lag) and how the middle class generates political pressure for liberalisation (some function of its size and relative income). Such a microfoundation would yield β_1 as a function of more primitive quantities and would allow the term to be disaggregated in countries where the middle-class mechanism does not operate (resource-rich autocracies, for instance).

6.4 Kramers escape time as a quantitative prediction

For each fitted country and each deadlock basin in its trajectory, the model implies a predicted mean first-passage time out of that basin. Comparing predicted to actual transition times across the panel would be a sharp test. We expect the model to perform poorly on transitions driven by mechanisms it does not represent (coups d'état driven by elite splits, decolonisation, foreign imposition), and well on transitions driven by gradual social change (modernisation-driven transitions).

6.5 Multi-regime extensions

The single-particle formulation isolates each polity from the others. A natural extension models several polities simultaneously as particles that exert forces on one another through $\boldsymbol{F}_{\text{ext}}$: the sanctioning state pushes the sanctioned state leftward; the liberalising bloc pulls neighbouring states rightward. This converts the model into a many-body system whose collective dynamics include cold-war-style polarisation, contagion of revolutions, and the formation of stable blocs. The empirical handle would be data on bilateral influence: trade exposure, sanctions, alliance structure.

6.6 Refining the state-capacity axis

State capacity is the dimension we are least confident about. The literature distinguishes at least three sub-components: fiscal, administrative, and coercive. A four-dimensional model with z split into z_f and z_c may yield a substantially better fit to cases such as oil-rich autocracies, which can have high z_f (revenue from extraction) and low z_c (limited bureaucratic reach into society). Whether the gain in fit justifies the loss in tractability is an open question.

6.7 Information environment as a control parameter

The current model has no explicit representation of the information environment, although in practice the cost of suppressing dissent (the $k_2 \exp(\cdot)$ term in Γ) is heavily mediated by information technology. A natural extension treats the cost of damping as a function of an exogenous technological parameter that has been falling over the last several decades. Such a treatment would formalise the widely-held intuition that information technology has shifted the deadlock landscape.

7 Conclusion

We have proposed a three-dimensional Langevin model of regime evolution that addresses the main defects of the earlier informal formulation. The model integrates the qualitative insights of modernisation theory, the state-capacity literature, and the political-economic theory of authoritarian persistence into a single tractable dynamical system, and it reproduces three distinct historical trajectories under a shared parameter set. The work that remains—empirical fitting, bifurcation analysis, microfoundations, multi-regime extension—is substantial. But the framework is now sharp enough that this work can be done, and the results can in principle be compared with something other than themselves. That is what we mean by "a model".

References

[1] F. Fukuyama, *Political Order and Political Decay*, Farrar, Straus and Giroux, 2014.

[2] P. Hänggi, P. Talkner, and M. Borkovec, "Reaction-rate theory: fifty years after Kramers," *Reviews of Modern Physics*, vol. 62, no. 2, pp. 251–341, 1990.

[3] J. K. Hanson and R. Sigman, "Leviathan's latent dimensions: measuring state capacity for comparative political research," working paper, 2017.

[4] M. Mann, "The autonomous power of the state: its origins, mechanisms and results," *Archives Européennes de Sociologie*, vol. 25, no. 2, pp. 185–213, 1984.

[5] M. Pei, *China's Trapped Transition: The Limits of Developmental Autocracy*, Harvard University Press, 2006.

[6] V-Dem Institute, "Varieties of Democracy dataset," University of Gothenburg.