Epigenetic Teleonomy: A Stochastic Control Model of Environmental Signaling

Elias Rubenstein *

Posted Date: 8 January 2026

doi: 10.20944/preprints202601.0583.v1

Keywords: epigenetic regulation; stochastic control; mean reversion; allostasis; divergence-to-baseline; relative entropy; Ornstein–Uhlenbeck process; Euler–Maruyama simulation

 doi:10.20944/preprints202601.0583.v1

Article

Epigenetic Teleonomy: A Stochastic Control Model of Environmental Signaling

Elias Rubenstein

Independent Researcher, USA; science@elias-rubenstein.com

Abstract

Background: Epigenetic regulation must preserve stable functional states under molecular stochasticity and changing environments, yet an operational model linking context-level signals to measurable chromatin remodeling is limited. **Method:** This study proposes *Epigenetic Teleonomy*, a stochastic control framework in which epigenomic observables relax toward an empirically estimated within-subject baseline regime (reference distribution) with lagged mediator-driven inputs and feedback. **Results:** A local approximation yields mean-reverting dynamics. Simulations illustrate that without effective feedback, diffusion-like drift leads to increasing dispersion, whereas sufficient regulation gain yields bounded fluctuations and recovery. In the isotropic local Ornstein–Uhlenbeck (OU) regime, stationary fluctuations scale with effective diffusion and inversely with return rate (gain). **Conclusion:** The framework is testable in longitudinal designs by (i) estimating a subject-specific baseline from a stable run-in window, (ii) quantifying deviation using reduced-dimensional proxies, and (iii) fitting gain and diffusion from return-to-baseline statistics.

Keywords: epigenetic regulation; stochastic control; mean reversion; allostasis; divergence-to-baseline; relative entropy; Ornstein–Uhlenbeck process; Euler–Maruyama simulation

1. Introduction

Epigenetics has matured from static "histone code" metaphors into an experimentally grounded account of chromatin state, memory, and regulatory plasticity [1,2]. Yet a persistent explanatory gap remains: how do macroscopic variables—such as environmental context, learning history, or structured inputs—translate into measurable molecular reorganizations with module-level specificity?

Mechanistic accounts explain how marks persist once set (maintenance), but often leave the directionality of reorganization implicit. Moreover, chromatin dynamics are continually perturbed by biochemical stochasticity, thermal fluctuations, and fluctuating cellular resources. In an open, non-equilibrium biological system, maintaining a stable functional regime is expected to rely on ongoing dissipative regulation and feedback: absent effective feedback, diffusion-like drift in state space can dominate. This motivates an explicit stochastic control model of regulation under perturbations; information-theoretic notions are used operationally (via divergence-to-baseline) rather than as a thermodynamic claim [6,8].

This paper proposes that chromatin regulation can be formalized as a *reference-regime control problem*. The concept of **Epigenetic Teleonomy** is introduced: the hypothesis that stability is maintained by feedback-driven minimization of deviation from an operational reference regime. Here, "teleonomy" is used only as a label for present-time error-correction (feedback control) toward an empirically defined reference regime, consistent with allostatic regulation [3].

A stochastic control model is formulated to replace vague metaphors with explicit equations governing the relaxation of epigenetic observables toward a target distribution. Stability conditions are derived in a local approximation and illustrated through numerical simulation.

 doi:10.20944/preprints202601.0583.v1

Contributions

This paper contributes: (i) a reference-regime stochastic control formulation of epigenetic stability with explicit gain/lag/coupling parameters; (ii) a local stability analysis that links return-to-reference rates and stationary fluctuation levels to identifiable quantities; and (iii) an operational falsification pathway for longitudinal designs using reduced-dimensional divergence proxies.

Intended Empirical Setting (Assay-Level Anchoring)

The intended use case is longitudinal, within-subject epigenomic monitoring in a fixed assay space (e.g., a fixed CpG panel, fixed peak set, or a fixed latent embedding learned on baseline samples). Measurements are taken repeatedly over time (days to weeks) with a stable run-in segment used to estimate the subject-specific reference regime. Context mediators (e.g., cortisol/cytokines) are measured on the same timeline to operationalize the lagged input term in Eq. (3).

Positioning relative to generic state-space models

While mean-reverting state-space models are well known, the distinguishing feature here is the explicit reference-regime formulation: a subject-specific baseline distribution Q, a divergence-defined potential $\Phi(\theta) = D_{KL}(P_y(\cdot|\theta)\|Q)$, and an explicit lagged actuation structure $(\mathbf{H} * (\mathbf{B}_s\mathbf{s}))(t)$ that connects measurable mediators to remodeling axes. The choice $\Phi(\theta) = D_{KL}(\cdot\|Q)$ is an *operational control objective* (a modeling choice) rather than a quantity derived from biochemistry; alternative deviation functionals can be substituted if empirically better supported. This yields two directly testable targets from time series: an estimated return rate and an estimated stationary variance; in the local isotropic OU approximation, stationary variance is expected to decrease as the estimated return rate increases under comparable noise conditions.

2. Formalizing the Epigenetic State Space

To bridge macroscopic signals and microscopic states, a formalism compatible with Waddington's landscape [10] but stated in operational, information-theoretic terms is adopted.

2.1. Microstate and Observable Distributions

Let $\mathbf{x} \in \mathcal{X}$ denote a high-dimensional microstate vector (e.g., methylation states at Cytosine-phosphate-Guanine [CpG] sites, chromatin accessibility states, or combinatorial mark configurations). The organism's instantaneous epigenetic condition is represented as a probability distribution $P(\mathbf{x}, t)$.

Experiments observe projections of this state (e.g., Assay for Transposase-Accessible Chromatin [ATAC-seq] peak counts, methylation beta values, mark enrichments). Let $\mathbf{y} \in \mathcal{Y}$ denote these observables and $P_y(\mathbf{y}, t)$ the corresponding distribution. For modeling, a low-dimensional parameterization is assumed:

$$P_y(\cdot, t) \approx P_y(\cdot \mid \theta(t)), \tag{1}$$

where $\theta(t)$ collects order parameters (e.g., module means/covariances, latent factors, mixture weights).

2.2. Operationalization (How θ, P_y, and Q are Defined from Data)

To keep the reference regime empirically grounded, the framework assumes:

- **Choice of observables:** $\mathbf{y}$ corresponds to a fixed assay space (e.g., a fixed CpG panel or peak set) observed across time.
- **Low-dimensional embedding:** $\theta(t)$ can be defined by a latent model (e.g., PCA/FA on standardized features, a Gaussian state-space model, or mixture components fitted to P_y at each time point).
- **Reference regime:** $Q(\cdot)$ is estimated from within-subject baselines during a stable run-in window (e.g., pooled baseline samples), rather than by population priors, to avoid circular definitions of "health."

 doi:10.20944/preprints202601.0583.v1

Parameter identification can proceed via state-space inference (Kalman-type filters for linear-Gaussian models or particle methods for nonlinear cases), with κ and Σ estimated from longitudinal variability and return-to-baseline dynamics.

2.3. The Reference Attractor $Q(\mathbf{y})$

A central feature is the definition of a stable reference regime, denoted $Q(\mathbf{y})$. Deviation from this reference is quantified by the Kullback–Leibler divergence (relative entropy):

$$D(t) \equiv D_{KL}(P_y(\cdot,t)\|Q) \;=\; \int_{\mathcal{Y}} P_y(\mathbf{y},t)\,\ln\frac{P_y(\mathbf{y},t)}{Q(\mathbf{y})}\,d\mathbf{y}. \tag{2}$$

This divergence is well-defined (finite) when $P_y(\cdot,t)$ is absolutely continuous with respect to Q; otherwise D_{KL} is infinite. Importantly, D_{KL} is *not* the thermodynamic entropy of the system; it is a measure of distributional deviation from the chosen reference.

3. Dynamics: The Control Loop

3.1. Controlled Diffusion Model

The evolution of $\theta(t)$ is modeled as a controlled diffusion that relaxes toward the reference basin defined by Q:

$$d\theta(t) = -\kappa\,\nabla_\theta\Phi(\theta(t))\,dt \;+\; \big(\mathbf{H}*(\mathbf{B}_s\,\mathbf{s})\big)(t)\,dt \;+\; \Sigma^{1/2}\,d\mathbf{W}(t), \tag{3}$$

where:

- $\Phi(\theta) \equiv D_{KL}(P_y(\cdot\mid\theta)\|Q)$ is a divergence-defined potential (a distance-to-reference objective). This identification is an *operational modeling choice* (control objective), not a biochemical derivation.
- $\kappa \geq 0$ is the **regulation gain** (aggregate strength of error-correction).
- $\mathbf{s}(t)$ is a measurable mediator vector (e.g., cortisol, cytokines) representing transduced context.
- $\mathbf{B}_s$ is the actuation map linking mediators to epigenetic parameters.
- $\mathbf{H}(\tau)$ is a lag/hysteresis kernel; the convolution is $(\mathbf{H}*(\mathbf{B}_s\,\mathbf{s}))(t) = \int_0^\infty \mathbf{H}(\tau)\,\mathbf{B}_s\,\mathbf{s}(t-\tau)\,d\tau$.
- $\Sigma^{1/2}d\mathbf{W}(t)$ is stochastic noise (Wiener process).

We assume $P_y(\cdot\mid\theta)$ is a smooth parametric family with common support on the region where $Q(\mathbf{y}) > 0$, so that $\Phi(\theta) = D_{KL}(P_y(\cdot\mid\theta)\|Q)$ is finite and differentiable and $\nabla_\theta\Phi(\theta)$ is well-defined.

3.2. Minimal Test Protocol (Operational Falsification)

Let a subject-specific run-in window define the reference regime Q in a fixed feature space or latent embedding. In practice, direct estimation of $D_{KL}(P_y\|Q)$ from finite samples in high dimension is often ill-posed; therefore define $\Delta(t)$ as a reduced-dimensional proxy for $D(t)$ (Eq. (2)) (e.g., Gaussian KL in latent space, MMD, or Wasserstein distance). Fit a baseline null model as a mean-reverting state-space model without mediator-driven input (set $\mathbf{B}_s = \mathbf{0}$). Evidence for structured steering requires that mediator terms improve out-of-sample prediction of $\Delta(t)$ and recovery timing beyond the null, and that estimated stationary fluctuation levels decrease as estimated return rates increase under comparable noise conditions.

3.3. Local Quadratic Approximation and Stability Intuition

Near a stable reference basin, a second-order Taylor expansion yields:

$$\Phi(\theta) \approx \frac{1}{2}\,(\theta-\theta^*)^\top A\,(\theta-\theta^*), \qquad A = \nabla^2\Phi(\theta^*) \succeq 0, \tag{4}$$

so $\nabla\Phi(\theta) \approx A(\theta-\theta^*)$. With bounded inputs and A positive definite in the local basin, the feedback term produces return-to-reference behavior of Ornstein–Uhlenbeck type [11].

 doi:10.20944/preprints202601.0583.v1

4. Numerical Simulation and Results

To illustrate the interplay between diffusion and feedback in the local regime, the linearized case with $A = I$ and $\theta^* = \mathbf{0}$ is simulated:

$$d\theta(t) = -\kappa\theta(t)\,dt + \sigma\,d\mathbf{W}(t), \tag{5}$$

which is the multivariate Ornstein–Uhlenbeck process. This is an *illustration* of local return-to-reference dynamics, not a claim that global epigenetic regulation is linear. In this simulation, the noise structure corresponds to the isotropic choice $\Sigma = \sigma^2 I$ in Eq. (3).

4.1. Simulation Setup

- **Dimension:** $N = 2$ for visualization.
- **Initial condition:** $\theta_0 = [2.0, 1.5]$.
- **Integration:** Euler–Maruyama scheme [12,13].
- **Fair comparison:** The regulated and unregulated trajectories use the same Wiener increments to isolate the effect of the drift term.

4.2. Results

Figure 1 shows:

1. **Drift vs. control:** With $\kappa = 0$, trajectories behave as Brownian motion with variance that grows in time.
2. **Bounded fluctuations:** With $\kappa > 0$, trajectories are constrained to a bounded region around the reference.
3. **Return-to-reference proxy:** In the OU regime, a natural local proxy for distance-to-reference is $\|\theta(t)\|^2$, which decays in expectation toward a stationary fluctuation level under feedback.

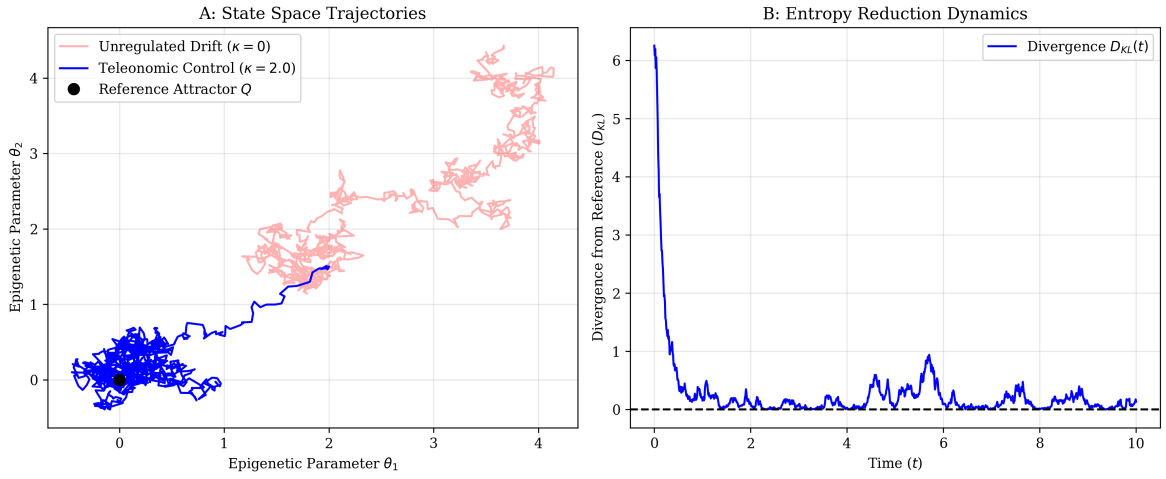

Figure 1. Local return-to-reference dynamics in the linearized regime. (A) State space trajectories for identical noise realizations: without feedback (red, $\kappa = 0$) the path diffuses, while feedback (blue, $\kappa = 2$) yields return-to-reference dynamics. **(B)** Local quadratic distance-to-reference proxy $\|\theta(t)\|^2$ in the OU approximation: the regulated regime exhibits decay (in expectation) toward a stationary fluctuation level, consistent with bounded variance scaling with diffusion strength and inverse gain in the isotropic case.

5. Discussion

5.1. Model Limitations and Scope

The OU dynamics used in simulation are a local approximation. Far-from-equilibrium transitions and multi-attractor landscapes require nonlinear $\Phi(\theta)$ and potentially state-dependent diffusion. The

 doi:10.20944/preprints202601.0583.v1

present contribution is therefore best interpreted as a *control-theoretic template* for connecting measured signals to return-to-reference dynamics, with explicit parameters that can be estimated.

5.2. Biological Interpretation of Parameters

The regulation gain κ corresponds to aggregate chromatin maintenance and remodeling activity. The coupling $\mathbf{B}_s$ captures how mediators project onto sensitive regulatory axes. Specific molecular substrates include ATP-dependent chromatin remodelers (e.g., SWI/SNF complex) and modifying enzymes (e.g., DNMTs, HATs); mechanistic background is reviewed in [5]. The diffusion term Σ captures effective stochasticity from biochemical fluctuations and resource variability.

5.3. Diffusion Bridges and a Constrained-Process Analogy (Teleonomy as Present-Time Drift)

A common mathematical representation of "goal constraints" in stochastic systems is *conditioning* a diffusion on an endpoint event (a bridge process). Consider an Itô diffusion

$$dX_t = b(t, X_t)\, dt + \sigma(t, X_t)\, dW_t, \qquad a = \sigma\sigma^\top. \tag{6}$$

Conditioning on an endpoint (e.g., $X_T = x_T$) yields a bridge whose drift can be written via Doob's h-transform:

$$dX_t = \Big(b(t, X_t) + a(t, X_t)\nabla_x \ln h(t, X_t)\Big)\, dt + \sigma(t, X_t)\, dW_t, \tag{7}$$

with $h(t, x)$ proportional to a transition density toward the conditioned event [14,15]. For Brownian motion this recovers the classical Brownian bridge drift. In our setting, this provides an *analogy*: global viability constraints can be represented as additional present-time drift terms. In local regimes, such drifts can be approximated by linear restoring forces, motivating why a feedback form like $-\kappa\nabla\Phi$ is a natural modeling choice. This is an analogy rather than an identity: the exact correspondence depends on the specific conditioning event and the associated h.

5.4. Implications for Intervention Specificity (Projection Condition)

The model clarifies why not every input produces specific remodeling: effective steering requires that the lagged signal $(\mathbf{H} * (\mathbf{B}_s\mathbf{s}))(t)$ has a non-zero projection onto sensitive directions of the local Hessian landscape. If inputs are orthogonal to these directions, return-to-reference behavior is governed mainly by feedback and diffusion rather than by signal-induced steering.

6. Conclusions

Epigenetic Teleonomy provides a falsifiable framework: effective regulation should manifest as a measurable reduction in divergence-to-reference relative to an empirically defined baseline, while stability requires sufficient gain relative to diffusion strength. Future work should estimate κ, Σ, $\mathbf{B}_s$, and $\mathbf{H}$ from time-series epigenomic measurements.

Author Contributions: The author conceived the theoretical framework, performed the mathematical derivations, implemented the numerical simulations, and wrote the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: In this section, you should add the Institutional Review Board Statement and approval number, if relevant to your study. You might choose to exclude this statement if the study did not require ethical approval. Please note that the Editorial Office might ask you for further information. Please add "The study was conducted in accordance with the Declaration of Helsinki, and approved by the Institutional Review Board (or Ethics Committee) of NAME OF INSTITUTE (protocol code XXX and date of approval)." for studies involving humans. OR "The animal study protocol was approved by the Institutional Review Board (or Ethics Committee) of NAME OF INSTITUTE (protocol code XXX and date of approval)." for studies involving animals. OR "Ethical review and approval were waived for this study due to REASON (please provide a detailed justification)." OR "Not applicable" for studies not involving humans or animals.

 doi:10.20944/preprints202601.0583.v1

Informed Consent Statement: Not applicable.

Data Availability Statement: No external datasets were used. All results are reproducible from the simulation code provided in Appendix B, which generates Figure 1.

Conflicts of Interest: The author declares no conflict of interest.

Abbreviations and Notation

See Table 1 for abbreviations and Table 2 for notation.

Table 1. Abbreviations used in this manuscript.

Abbreviation	Definition
SDE	Stochastic differential equation
OU	Ornstein–Uhlenbeck (return-to-reference) process
CpG	Cytosine-phosphate-Guanine dinucleotide
DNMT	DNA methyltransferase
HAT	Histone acetyltransferase
SWI/SNF	Switch/Sucrose Non-Fermentable (chromatin remodeling complex)
ATAC-seq	Assay for Transposase-Accessible Chromatin with high-throughput sequencing

Table 2. Notation used in this manuscript.

Symbol	Description
$\mathbf{x} \in \mathcal{X}$	High-dimensional epigenetic microstate vector
$P(\mathbf{x}, t)$	Instantaneous probability distribution of the microstate
$\mathbf{y} \in \mathcal{Y}$	Observable assay vector (projection of microstate)
$P_y(\mathbf{y}, t)$	Observable distribution parametrized by θ
$\theta(t)$	Epigenetic state vector (order parameters)
$Q(\mathbf{y})$	Reference regime distribution (within-subject baseline)
D_{KL} / $D(t)$	Kullback–Leibler divergence (relative entropy to reference)
$\Phi(\theta)$	Potential defined as $\Phi(\theta) = D_{KL}(P_y(\cdot\|\theta)\|Q)$
κ	Regulation gain (feedback strength)
$\mathbf{s}(t)$	Context signal vector (mediators)
$\mathbf{B}_s$	Actuation/coupling matrix
$\mathbf{H}(\tau)$	Temporal kernel (lag/hysteresis)
Σ	Noise covariance matrix
$\mathbf{W}(t)$	Wiener process (Brownian motion)
A	Local Hessian of Φ at θ^*
$\mathbb{E}[\cdot]$	Expectation operator

Appendix A Stability in the Local Quadratic Regime

Appendix A.1 A.1 Local quadratic approximation

Assume θ^* is a local reference point (e.g., a mode of the reference basin) and

$$\Phi(\theta) \approx \frac{1}{2}(\theta - \theta^*)^\top A(\theta - \theta^*), \qquad A = \nabla^2\Phi(\theta^*) \succeq 0.$$

Then $\nabla\Phi(\theta) \approx A(\theta - \theta^*)$ and, neglecting bounded inputs for the moment, Eq. (3) becomes

$$d\theta = -\kappa\, A(\theta - \theta^*)\, dt + \Sigma^{1/2}\, d\mathbf{W}.$$

If A is positive definite in the basin, the process is mean-reverting (OU-type).

 doi:10.20944/preprints202601.0583.v1

Appendix A.2 A.2 Lyapunov drift and stationary scaling

Let $V(\theta) = \frac{1}{2}\|\theta - \theta^*\|^2$. By Itô's Lemma [13],

$$dV = (\nabla V)^\top d\theta + \frac{1}{2}\,\mathrm{Tr}(\nabla^2 V\,\Sigma)\,dt.$$

With $\nabla V = \theta - \theta^*$ and $\nabla^2 V = I$:

$$\frac{d}{dt}\mathbb{E}[V] = -\kappa\,\mathbb{E}\big[(\theta-\theta^*)^\top A(\theta-\theta^*)\big] + \frac{1}{2}\,\mathrm{Tr}(\Sigma).$$

For the linearized multivariate OU system, the stationary covariance C satisfies the continuous-time Lyapunov equation

$$\kappa\,(AC + CA^\top) = \Sigma.$$

In the isotropic special case $A = I$, this reduces to $C = \Sigma/(2\kappa)$ (and for scalar θ gives $\mathrm{Var}(\theta) = \Sigma/(2\kappa)$).

Appendix A.3 A.3 Bounded-input remark (input-to-state stability intuition)

If the lagged input term $(\mathbf{H} * (\mathbf{B}_s\mathbf{s}))(t)$ is bounded in norm and $A \succ 0$ locally, standard linear systems intuition implies the state remains bounded in expectation around a shifted mean. This motivates treating signal transduction as steering within (or between) basins while feedback stabilizes against diffusion.

Appendix A.4 A.4 Diffusion bridges and Doob's h-transform (formal statement)

For an Itô diffusion $dX_t = b(t, X_t)\,dt + \sigma(t, X_t)\,dW_t$ with $a = \sigma\sigma^\top$, conditioning on an endpoint event can be represented by a Doob h-transform, yielding an additional drift term $a\nabla \log h$ under suitable regularity assumptions [14,15]. This provides a principled way to interpret constraints as present-time drift terms, though the exact mapping depends on the chosen conditioning event and transition kernel.

Appendix B Appendix B: Python Code for Numerical Simulation

Listing 1: Reproducible simulation and plotting (OU illustration)

```
import numpy as np
import matplotlib.pyplot as plt

# Reproducibility
np.random.seed(42)

# 1. Parameters
T = 10.0
dt = 0.01
N = int(T / dt) + 1
time = np.linspace(0.0, T, N)

kappa = 2.0             # regulation gain
sigma = 0.5             # diffusion strength (scalar for simplicity)

# 2. State initialization (2D)
theta_reg = np.zeros((N, 2))
theta_drift = np.zeros((N, 2))

start_pos = np.array([2.0, 1.5])
```

 doi:10.20944/preprints202601.0583.v1


```
theta_reg[0] = start_pos
theta_drift[0] = start_pos

# 3. Pre-generate Wiener increments for fair comparison
dW_all = np.random.normal(0.0, np.sqrt(dt), size=(N-1, 2))

# 4. Euler-Maruyama integration
for t in range(N - 1):
    dW = dW_all[t]

    # Regulated OU: dtheta = -kappa * theta * dt + sigma * dW
    theta_reg[t + 1] = theta_reg[t] + (-kappa * theta_reg[t]) * dt + sigma * dW

    # Unregulated drift: dtheta = sigma * dW
    theta_drift[t + 1] = theta_drift[t] + sigma * dW

# 5. Local distance-to-reference proxy (quadratic distance)
dist_reg = np.sum(theta_reg**2, axis=1)
dist_drift = np.sum(theta_drift**2, axis=1)

# 6. Plot: combined figure (A) phase plot (B) distance proxy
fig = plt.figure(figsize=(12, 5))

# (A) State space trajectories
ax1 = fig.add_subplot(1, 2, 1)
ax1.plot(theta_drift[:, 0], theta_drift[:, 1], label=r"Unregulated (\kappa=0)")
ax1.plot(theta_reg[:, 0], theta_reg[:, 1], label=r"Regulated (\kappa=2)")
ax1.scatter([0], [0], marker="x", s=60, label=r"Reference (\theta^\ast=0)")
ax1.set_xlabel(r"\theta_1")
ax1.set_ylabel(r"\theta_2")
ax1.set_title("State space trajectories")
ax1.legend(frameon=False)

# (B) Quadratic distance proxy over time
ax2 = fig.add_subplot(1, 2, 2)
ax2.plot(time, dist_drift, label=r"Unregulated (\kappa=0)")
ax2.plot(time, dist_reg, label=r"Regulated (\kappa=2)")
ax2.set_xlabel("time")
ax2.set_ylabel(r"\|\theta(t)\|^2 (local proxy)")
ax2.set_title("Local distance-to-reference proxy")
ax2.legend(frameon=False)

plt.tight_layout()
plt.savefig("sim_results_combined.png", dpi=300)
plt.close(fig)
```

References

1. Bird, A. (2007). Perceptions of epigenetics. *Nature*, 447, 396–398. doi:10.1038/nature05913
2. Jenuwein, T., & Allis, C. D. (2001). Translating the histone code. *Science*, 293, 1074–1080. doi:10.1126/science.1063125

 doi:10.20944/preprints202601.0583.v1

3. McEwen, B. S., & Wingfield, J. C. (2003). The concept of allostasis in biology and biomedicine. *Hormones and Behavior*, 43, 2–15. doi:10.1016/S0018-506X(02)00024-5
4. Teschendorff, A. E., et al. (2016). DNA methylation outliers in normal breast tissue identify field defects that are enriched in cancer. *Nature Communications*, 7, 10478. doi:10.1038/ncomms10478
5. Clapier, C. R., et al. (2017). Mechanisms of action and regulation of ATP-dependent chromatin-remodelling complexes. *Nature Reviews Molecular Cell Biology*, 18, 407–422. doi:10.1038/nrm.2017.26
6. Landauer, R. (1961). Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development*, 5, 183–191. doi:10.1147/rd.53.0183
7. Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27, 379–423. doi:10.1002/j.1538-7305.1948.tb01338.x
8. Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138. doi:10.1038/nrn2787
9. Levin, M. (2019). The Computational Boundary of a "Self". *Frontiers in Psychology*, 10, 2688. doi:10.3389/fpsyg.2019.02688
10. Waddington, C. H. (1957). *The Strategy of the Genes*. Allen & Unwin. (Reprint 2014 Routledge, ISBN: 9781138017313).
11. Uhlenbeck, G. E., & Ornstein, L. S. (1930). On the Theory of the Brownian Motion. *Physical Review*, 36, 823. doi:10.1103/PhysRev.36.823
12. Kloeden, P. E., & Platen, E. (1992). *Numerical Solution of Stochastic Differential Equations*. Springer. doi:10.1007/978-3-662-12616-5
13. Øksendal, B. (2003). *Stochastic Differential Equations: An Introduction with Applications* (6th ed.). Springer. doi:10.1007/978-3-642-14394-6
14. Doob, J. L. (1957). Conditional Brownian motion and the boundary limits of harmonic functions. *Bulletin de la Société Mathématique de France*, 85, 431–458. doi:10.24033/bsmf.1494
15. Baudoin, F. (2002). Conditioned stochastic differential equations: Theory, examples and application to finance. *Stochastic Processes and their Applications*, 100(1–2), 109–145. doi:10.1016/S0304-4149(02)00109-6