Arithmetic-Sieve Random Differential Regime Models

The Zeta–Möbius RDE Regime Probability Engine

P. Zavari

August 2026

Abstract

This paper develops a current-conditioned regime probability engine that combines a causal Zeta–Möbius arithmetic-sieve transform with polynomially forced Random Differential Equations (RDEs). The construction begins from Euler's product for the Riemann zeta function and the reciprocal Dirichlet expansion in the Möbius function. The resulting finite transform is used as a signed structural-memory state for financial time series. That state is embedded with log price and log volatility in a three-coordinate RDE system whose deterministic forecast trajectories are obtained from the closed Wiener-forcing expectation formula of the structured-forcing RDE framework. We derive the Wiener resolvent kernels, prove an exact endpoint-conditioning identity, establish the weighted ridge calibration used by the engine, construct model-evidence scenario weights and regime-cluster probabilities, formalize downside and volatility outputs, and derive a robust historical-manifold anomaly statistic. We also prove that a fixed set of continuous trajectories yields continuous scenario-band envelopes, separating legitimate regime jumps from interval-selection artifacts, and present a hybrid jump-RDE extension in which vertical state discontinuities correspond to explicit regime transitions. The resulting framework reports regime classification, scenario likelihoods, expected returns, downside quantiles, volatility paths, Zeta–Möbius transitions, trajectory disagreement, and distance from the historical market manifold. The scenario probabilities are model-evidence weights over deterministic specifications; they are not claimed to be risk-neutral probabilities or empirically calibrated posterior probabilities without additional validation.

Keywords: random differential equations; Möbius function; Riemann zeta function; structured stochastic forcing; regime detection; scenario probabilities; market manifold; volatility; risk management.

Contents

1 Introduction

Financial regime models face two distinct problems. The first is *state construction*: what quantities should represent the current organization of the market? The second is *state evolution*: once a state has been defined, how should it be projected forward without replacing all nonlinear structure by either a linear time-series model or Monte Carlo noise?

The framework developed here addresses these two problems with different mathematical objects. The Zeta–Möbius transform supplies a causal arithmetic-sieve state. It is deterministic, signed, sparse over lags, and derived from the reciprocal of the Riemann zeta function. The RDE supplies the dynamics. Rather than injecting randomness through an infinitesimal stochastic integral, the RDE treats a stochastic path as an external forcing signal and solves an ordinary differential equation pathwise. This allows polynomial transformations of Wiener paths to be handled by ordinary resolvent calculus while retaining explicit expectation formulas [7, 4, 5].

The practical engine uses the joint state

$$X_t = \big(L_t,\, Z_t,\, V_t\big)^\top, \tag{1}$$

where $L_t = \log P_t$ is log adjusted price, Z_t is a causal Zeta–Möbius state computed from returns, and V_t is log volatility. Each coordinate is robustly standardized and fitted by a current-ending polynomial-forcing RDE. A grid of calibration windows, ridge penalties, and recency weights produces a deterministic ensemble of current-conditioned trajectories. The ensemble is converted into model-evidence weights, clustered into regimes, and summarized into the quantities shown in the dashboard.

The central mathematical contributions assembled in this paper are:

1. a precise causal Zeta–Möbius operator and its boundedness and truncation properties;
2. the closed Wiener-forcing RDE expectation and resolvent kernels;
3. an exact current-endpoint conditioning theorem that forces every projected trajectory through today's observed state without changing the fitted forcing law;
4. a weighted ridge formulation for fitting the effective polynomial forcing coefficients;
5. deterministic model-evidence weights and regime probabilities;
6. coherent pathwise uncertainty tubes for volatility and Zeta–Möbius trajectories;
7. a robust historical-manifold metric for anomaly and regime-change detection; and
8. a jump-RDE extension that makes regime shifts explicit rather than allowing plotting artifacts to masquerade as structural jumps.

The construction should be interpreted as a state-estimation and scenario engine, not as an arbitrage-free asset-pricing theory. In particular, when the engine plots $\exp(\mathbb{E}[L_t])$, that is a geometric central path; in general $\exp(\mathbb{E}[L_t]) \neq \mathbb{E}[\exp(L_t)]$ by Jensen's inequality. A complete conditional expectation of price would also require the conditional variance and higher distributional information of log price.

2 Number-theoretic source: the Zeta–Möbius sieve

2.1 Euler product and reciprocal expansion

For $\mathrm{Re}(s) > 1$, the Riemann zeta function is

$$\zeta(s) = \sum_{n=1}^{\infty} n^{-s}. \tag{2}$$

Euler's product expresses the same analytic object through prime factorization:

$$\zeta(s) = \prod_{p \in \mathcal{P}} \frac{1}{1 - p^{-s}}, \tag{3}$$

where $\mathcal{P}$ is the set of primes. Taking the reciprocal gives

$$\frac{1}{\zeta(s)} = \prod_{p \in \mathcal{P}} (1 - p^{-s}). \tag{4}$$

Expanding the product selects each prime either zero or one time. Terms containing a repeated prime factor are absent. This produces the Möbius function

$$\mu(n) = \begin{cases} 1, & n = 1, \\ (-1)^k, & n = p_1 \cdots p_k \text{ with distinct primes } p_j, \\ 0, & p^2 \mid n \text{ for some prime } p. \end{cases} \tag{5}$$

Thus

$$\frac{1}{\zeta(s)} = \sum_{n=1}^{\infty} \frac{\mu(n)}{n^s}, \qquad \mathrm{Re}(s) > 1. \tag{6}$$

This is standard analytic number theory [1, 2]. The financial construction does not claim that market observations obey a number-theoretic generative law. Instead, it imports the sign/zero pattern of (6) as a deterministic cancellation kernel.

2.2 Finite causal transform

Let X_t be a discrete signal observed at integer time t. Examples include returns, volume, spread, order-flow imbalance, realized variance, funding rate, or implied-volatility skew.

Definition 2.1 (Causal Zeta–Möbius transform)**.** For truncation depth $N \in \mathbb{N}$ and exponent $s > 0$, define

$$Z_t^{(N,s)}[X] = \sum_{n=1}^{N} \frac{\mu(n)}{n^s} X_{t-n}. \tag{7}$$

A scale-normalized version is

$$\widetilde{Z}_t^{(N,s)}[X] = \sum_{n=1}^{N} \widetilde{w}_n^{(N,s)} X_{t-n}, \qquad \widetilde{w}_n^{(N,s)} = \frac{\mu(n) n^{-s}}{\sum_{k=1}^{N} |\mu(k)| k^{-s}}. \tag{8}$$

The production regime engine described later uses the raw finite kernel in (7); the normalized form is useful when comparing several values of N or s on a common scale. The previous market-microstructure formulation used normalized kernels as a multi-scale feature stack [8].

Theorem 2.2 (Basic properties of the finite arithmetic sieve)**.** *For any $N \geq 1$ and $s > 0$:*

(i) $Z_t^{(N,s)}$ is causal and linear in X.

(ii) Lags divisible by the square of a prime have exactly zero coefficient.

(iii) For the normalized transform,

$$|\widetilde{Z}_t^{(N,s)}[X]| \leq \max_{1 \leq n \leq N} |X_{t-n}|. \tag{9}$$

(iv) If $s > 1$ and $|X_t| \leq M$ for all t, then the infinite raw transform converges absolutely and the truncation error satisfies

$$\left| \sum_{n>N} \frac{\mu(n)}{n^s} X_{t-n} \right| \leq M \sum_{n>N} n^{-s} \leq \frac{M N^{1-s}}{s-1}. \tag{10}$$

Proof. Linearity and causality follow immediately from (7); all terms use times strictly before t. Property (ii) follows from $\mu(n) = 0$ whenever $p^2 \mid n$. For (iii), the triangle inequality and ℓ^1 normalization give

$$|\widetilde{Z}_t| \leq \sum_{n=1}^{N} |\widetilde{w}_n| \, |X_{t-n}| \leq \max_n |X_{t-n}| \sum_{n=1}^{N} |\widetilde{w}_n| = \max_n |X_{t-n}|.$$

For (iv), use $|\mu(n)| \leq 1$ and the integral test: $\sum_{n>N} n^{-s} \leq \int_N^\infty x^{-s}\,\mathrm{d}x = N^{1-s}/(s-1)$. □

2.3 Residual energy and structural interpretation

For a multi-scale collection $\mathcal{N} \times \mathcal{S}$ and channels $X^{(j)}$, define

$$\Phi_t = \left(Z_t^{(N,s)}[X^{(j)}] \right)_{N \in \mathcal{N},\, s \in \mathcal{S},\, j=1,\ldots,d}. \tag{11}$$

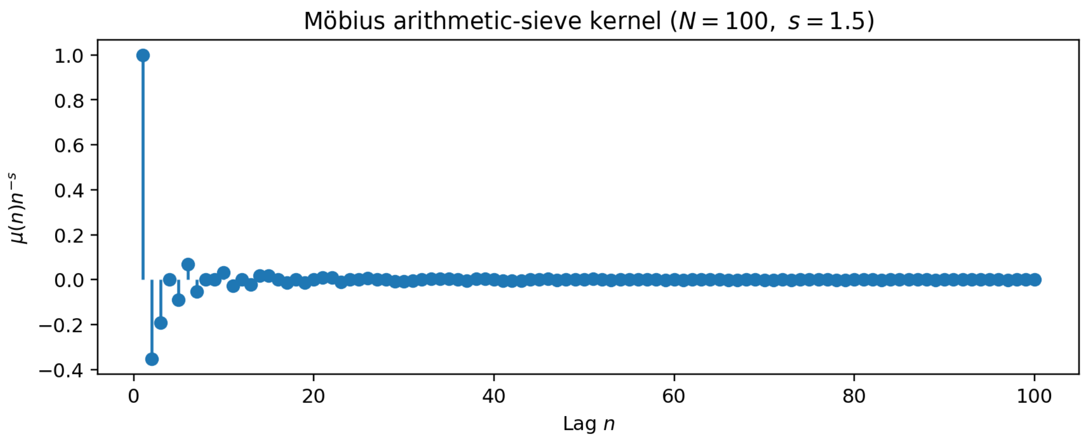

Figure 1: The raw arithmetic-sieve kernel $\mu(n)n^{-s}$ for $N = 100$ and $s = 1.5$. Square-containing lags vanish; square-free lags alternate sign according to the parity of the number of distinct prime factors.

A local arithmetic residual energy can be defined by

$$E_t = \sum_{j,N,s} \omega_{j,N,s} \operatorname{MA}_{h(N)} \left(\left[Z_t^{(N,s)}[X^{(j)}] \right]^2 \right), \qquad \omega_{j,N,s} \geq 0. \tag{12}$$

The interpretation is cancellation rather than literal randomness. Ordinary locally incoherent activity may cancel under the signed sieve, while a systematic lag organization may survive. A large $|Z_t|$ indicates that the recent history does not cancel under the chosen arithmetic lag pattern; a value near zero indicates stronger cancellation. The sign itself is an orientation of the signed lag balance and is not universally bullish or bearish.

3 Random differential equations with structured forcing

3.1 Pathwise RDE formulation

Consider the generalized Ornstein–Uhlenbeck RDE

$$\dot{x}(t) + A\big(x(t) - \mu\big) = F(LZ(t)), \qquad x(0) = x_0, \tag{13}$$

where $A \in \mathbb{R}^{d\times d}$ is stable, L maps the external process Z into correlated forcing coordinates, and F is polynomial or analytic. Unlike an Itô SDE, (13) contains the random path inside an ordinary time derivative; for each fixed realization of Z, it is an ODE [7].

Proposition 3.1 (Mild representation). *Assume A is constant and the forcing $G(t) = F(LZ(t))$ is locally integrable. Then any solution of* (13) *satisfies*

$$x(t) = \mu + e^{-At}(x_0 - \mu) + \int_0^t e^{-A(t-u)} G(u)\, \mathrm{d}u. \tag{14}$$

Proof. Set $y(t) = x(t) - \mu$. Then $\dot{y} + Ay = G(t)$. Multiply by e^{At} to obtain $\frac{\mathrm{d}}{\mathrm{d}t}(e^{At}y(t)) = e^{At}G(t)$. Integrating from 0 to t and multiplying by e^{-At} yields (14). □

If $A = P\Lambda P^{-1}$ is diagonalizable, the change of coordinates $\widetilde{x} = P^{-1}x$ reduces the linear resolvent to $\Lambda = \operatorname{diag}(\lambda_1, \dots, \lambda_d)$. The current engine works coordinatewise in such decoupled scalar modes.

4 Polynomial Wiener forcing and the Section 5.2 expectation

4.1 Coordinate representation

Assume $A = \operatorname{diag}(\lambda_1, \dots, \lambda_d)$ with $\lambda_i > 0$, $Z(t) = W(t)$ is standard Brownian motion, and

$$F(y) = \sum_{k=0}^{p} a_k y^{\odot k}. \tag{15}$$

Let

$$Y_i(t) = (LW(t))_i, \qquad \sigma_i^2 = \sum_m L_{im}^2. \tag{16}$$

Then

$$x_i(t) = \mu_i + e^{-\lambda_i t}(x_{0,i} - \mu_i) + \sum_{k=0}^{p} a_k \int_0^t e^{-\lambda_i (t-u)} Y_i(u)^k \,\mathrm{d}u. \tag{17}$$

Since $Y_i(u) \sim N(0, u\sigma_i^2)$,

$$\mathbb{E}[Y_i(u)^{2q+1}] = 0, \qquad \mathbb{E}[Y_i(u)^{2q}] = \frac{(2q)!}{2^q q!}\sigma_i^{2q} u^q. \tag{18}$$

Theorem 4.1 (Wiener polynomial-forcing expectation)**.** *Define*

$$J_q(t;\lambda) = \int_0^t e^{-\lambda(t-u)} u^q \,\mathrm{d}u. \tag{19}$$

Whenever expectation and integration may be interchanged,

$$\boxed{\mathbb{E}[x_i(t)] = \mu_i + e^{-\lambda_i t}(x_{0,i} - \mu_i) + \sum_{q=0}^{\lfloor p/2 \rfloor} a_{2q} \frac{(2q)!}{2^q q!} \sigma_i^{2q} J_q(t; \lambda_i).} \tag{20}$$

Proof. Take expectations in (17). By (18), all odd powers vanish. For $k = 2q$,

$$\mathbb{E}\left[\int_0^t e^{-\lambda_i (t-u)} Y_i(u)^{2q} \,\mathrm{d}u\right] = \frac{(2q)!}{2^q q!}\sigma_i^{2q} \int_0^t e^{-\lambda_i (t-u)} u^q \,\mathrm{d}u.$$

Substitution gives (20). □

This is the Wiener-input expectation derived in Section 5.2 of the structured-forcing RDE paper

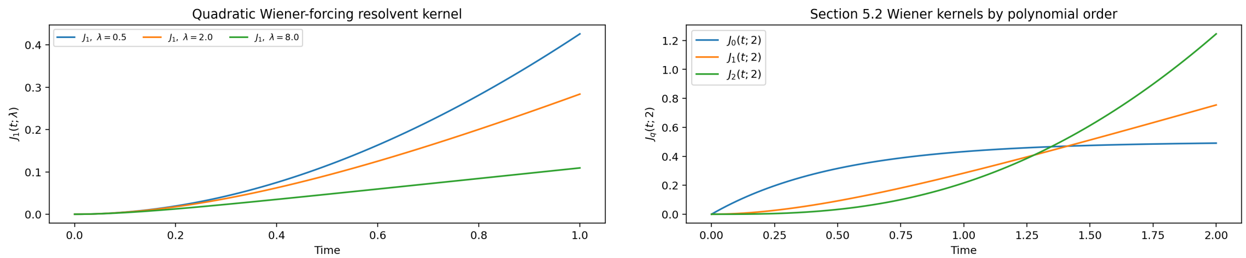

(a) Quadratic kernel across decay rates.

(b) Kernel growth across polynomial orders.

Figure 2: Section 5.2 Wiener resolvent kernels. Larger λ suppresses memory more rapidly; higher polynomial order creates stronger secular growth.

[7]. The coefficient

$$c_q = \frac{(2q)!}{2^q q!} = (2q-1)!! \tag{21}$$

is the even Gaussian moment factor.

4.2 Closed resolvent kernel

Proposition 4.2 (Recurrence and closed form of J_q)**.** *For $q \in \mathbb{N}_0$ and $\lambda > 0$,*

$$J_0(t;\lambda) = \frac{1-e^{-\lambda t}}{\lambda}, \tag{22}$$

$$J_q(t;\lambda) = \frac{t^q}{\lambda} - \frac{q}{\lambda} J_{q-1}(t;\lambda), \qquad q \geq 1, \tag{23}$$

and therefore

$$\boxed{J_q(t;\lambda) = \sum_{\beta=0}^{q} \frac{(-1)^\beta q!}{(q-\beta)!\lambda^{\beta+1}} t^{q-\beta} + \frac{(-1)^{q+1} q!}{\lambda^{q+1}} e^{-\lambda t}.} \tag{24}$$

Proof. For $q = 0$, integrate directly. For $q \geq 1$, write $J_q = e^{-\lambda t} \int_0^t e^{\lambda u} u^q \, du$ and integrate by parts with u^q and $e^{\lambda u}/\lambda$. This yields (23). Repeated substitution terminates at J_0 and gives (24). □

Two important special cases are

$$J_1(t;\lambda) = \frac{t}{\lambda} - \frac{1-e^{-\lambda t}}{\lambda^2}, \tag{25}$$

$$J_2(t;\lambda) = \frac{t^2}{\lambda} - \frac{2t}{\lambda^2} + \frac{2(1-e^{-\lambda t})}{\lambda^3}. \tag{26}$$

Corollary 4.3 (Secular growth)**.** *For fixed $\lambda > 0$,*

$$J_q(t;\lambda) = \frac{t^q}{\lambda} - \frac{q t^{q-1}}{\lambda^2} + O(t^{q-2}), \qquad t \to \infty, \tag{27}$$

with the obvious interpretation for $q = 0, 1$. In particular, nonzero quadratic Wiener forcing produces *a mean contribution growing linearly in time, while quartic forcing produces a leading quadratic contribution.*

Proof. Read the polynomial terms directly from (24); the final exponential term is $o(1)$. □

This is why the current production engine defaults to polynomial degree $p = 2$. Quartic terms can be valid on controlled finite horizons, but the J_2 term can dominate extrapolation rapidly.

4.3 Effective coefficient identifiability

The first moment depends on a_{2q} and σ_i only through

$$b_{i,2q} := a_{2q} \frac{(2q)!}{2^q q!} \sigma_i^{2q}. \tag{28}$$

Proposition 4.4 (First-moment non-identifiability)**.** *If only the mean trajectory (20) is observed, then a_{2q} and σ_i are not separately identified without additional restrictions; only $b_{i,2q}$ in (28) is identified.*

Proof. For any $c > 0$, replacing σ_i by $c\sigma_i$ and a_{2q} by $a_{2q}c^{-2q}$ leaves $b_{i,2q}$ unchanged for every q, and therefore leaves the mean trajectory unchanged. □

The regime engine therefore fits effective coefficients $b_{i,2q}$ directly.

5 Exact current-endpoint conditioning

Historical-window fitting alone is not sufficient for a current forecast. A fitted trajectory may miss today's observed state. For a linear resolvent, there is an exact correction.

Theorem 5.1 (Endpoint-conditioned resolvent)**.** *Let $m(t)$ solve*

$$m'(t) + \lambda\big(m(t) - \mu\big) = g(t), \qquad \lambda > 0. \tag{29}$$

Fix a conditioning time T and observed state x_T. For $t \geq T$, define

$$\boxed{m^{(T)}(t) = m(t) + e^{-\lambda(t-T)}\big[x_T - m(T)\big].} \tag{30}$$

Then

(i) $m^{(T)}(T) = x_T$;

(ii) $m^{(T)}$ solves the same forced ODE (29) for $t > T$;

(iii) it is the unique solution of (29) on $[T, \infty)$ with initial condition $m^{(T)}(T) = x_T$.

Proof. At $t = T$, the correction equals $x_T - m(T)$, proving (i). Differentiate (30):

$$(m^{(T)})'(t) = m'(t) - \lambda e^{-\lambda(t-T)}[x_T - m(T)].$$

Hence

$$\begin{aligned}(m^{(T)})'(t) + \lambda(m^{(T)}(t) - \mu) &= m'(t) + \lambda(m(t) - \mu) \\ &= g(t),\end{aligned}$$

which proves (ii). Uniqueness follows from the standard uniqueness theorem for the scalar linear ODE. □

Corollary 5.2 (Current derivative)**.** *If*

$$g(t) = \sum_q b_{2q} t^q, \tag{31}$$

then the right derivative of the conditioned path at T is

$$(m^{(T)})'(T^+) = -\lambda(x_T - \mu) + \sum_q b_{2q} T^q. \tag{32}$$

The formula in (30) is central to the dashboard. Every specification is fitted on a window that terminates today, and every projected state is forced through the same observed endpoint. This removes the earlier failure mode in which unrelated historical regimes were transplanted onto the current market state.

6 The three-coordinate market state

The engine models three market states:

$$R_t = \begin{bmatrix} L_t \\ Z_t \\ V_t \end{bmatrix} = \begin{bmatrix} \log P_t \\ Z_t^{(N,s)}[r] \\ \log \widehat{\sigma}_t \end{bmatrix}, \tag{33}$$

where $r_t = \log(P_t/P_{t-1})$ and $\widehat{\sigma}_t$ is an EWMA volatility estimate. The volatility state is stored in logarithmic units so positivity is automatic after inversion.

For robust calibration, each coordinate is transformed to

$$X_{t,i} = \frac{R_{t,i} - c_i}{s_i}, \tag{34}$$

where c_i is a robust center and $s_i > 0$ a robust scale, typically based on the median and MAD over a recent normalization window.

For the current degree-two engine, each coordinate is represented by

$$m_i(t) = e^{-\lambda_i t}x_{0,i} + (1 - e^{-\lambda_i t})\mu_i + b_{i,2}J_1(t;\lambda_i), \tag{35}$$

followed by current-endpoint conditioning (30).

The raw projected quantities are then

$$\widehat{L}_j(T+h) = c_L + s_L m_{j,L}^{(T)}(T+h), \tag{36}$$
$$\widehat{Z}_j(T+h) = c_Z + s_Z m_{j,Z}^{(T)}(T+h), \tag{37}$$
$$\widehat{V}_j(T+h) = c_V + s_V m_{j,V}^{(T)}(T+h), \tag{38}$$

and

$$\widehat{P}_j(T+h) = \exp\{\widehat{L}_j(T+h)\}, \tag{39}$$
$$\widehat{\sigma}_j^{\mathrm{ann}}(T+h) = \sqrt{D}\,\exp\{\widehat{V}_j(T+h)\}, \tag{40}$$

where D is the number of trading days per year.

Proposition 6.1 (Positivity)**.** *The projected price and volatility in* (39)–(40) *are strictly positive whenever the latent paths are finite.*

Proof. The exponential map takes every finite real number to $(0,\infty)$. □

Modeling log price directly is important. It avoids mechanically adding a positive historical return center to every future day, a construction that can produce an upward forecast bias even for a severely declining asset.

7 Current-ending calibration

7.1 Recency-weighted ridge fit

Fix a calibration window of M observations ending at the present. Let $t_\ell = \ell/D$, $\ell = 1,\dots,M$, and $T = M/D$. For fixed λ, write the Section 5.2 mean in linear form

$$y_\ell - e^{-\lambda t_\ell}x_0 = \beta_0(1 - e^{-\lambda t_\ell}) + \sum_{q=1}^{Q}\beta_q J_q(t_\ell;\lambda) + \varepsilon_\ell, \tag{41}$$

with $\beta_0 = \mu$ and $\beta_q = b_{2q}$. Recent observations receive larger weight

$$w_\ell = \exp\left[-\delta\frac{T - t_\ell}{T}\right], \qquad \delta \ge 0. \tag{42}$$

Thus $w_M = 1$ and the earliest observation has weight approximately $e^{-\delta}$.

Let X_λ denote the design matrix from (41), $W = \operatorname{diag}(w_1, \dots, w_M)$, and D_ρ a diagonal coefficient-penalty matrix.

Proposition 7.1 (Weighted ridge solution)**.** *For fixed λ, if $X_\lambda^\top W X_\lambda + \rho D_\rho$ is nonsingular, the unique minimizer of*

$$\|W^{1/2}(y - X_\lambda \beta)\|_2^2 + \rho \beta^\top D_\rho \beta \tag{43}$$

is

$$\widehat{\beta}(\lambda) = \left(X_\lambda^\top W X_\lambda + \rho D_\rho\right)^{-1} X_\lambda^\top W y. \tag{44}$$

Proof. Differentiate (43) with respect to β, set the gradient equal to zero, and solve the normal equations. □

A positive grid of λ values is searched and the value minimizing the penalized weighted error is selected. The complete ensemble varies window length, ridge strength, and recency decay. This is deterministic specification uncertainty, not Monte Carlo randomness.

7.2 Current-trend conditioning

Let g_i denote a robust estimate of the current derivative of state coordinate i, for example a weighted combination of slopes over several recent windows. For a fitted specification j, let $d_{j,i}$ be the conditioned derivative in (32). Define the scaled mismatch

$$M_j = \frac{\sum_i \omega_i |d_{j,i} - g_i| / a_i}{\sum_i \omega_i}, \tag{45}$$

where $a_i > 0$ prevents division by a near-zero trend and ω_i emphasizes economically important states. In the current implementation, log price receives the largest weight.

When the present log-price trend is sufficiently strong, candidate specifications may be required to satisfy

$$\operatorname{sgn}(d_{j,L}) = \operatorname{sgn}(g_L), \tag{46}$$

and, under the strict continuation setting,

$$\operatorname{sgn}\left(\widehat{L}_j(T+H) - L_T\right) = \operatorname{sgn}(g_L). \tag{47}$$

These are explicit modeling constraints, not theorems that past trends must continue. They are used when the purpose is to prevent a strongly declining current asset from being represented primarily by unrelated historical bull-regime fits.

7.3 Historical transition envelope

For the non-price state coordinates $i \in \{Z, V\}$ and horizon h, define the empirical transition envelope

$$E_i(h) = Q_\alpha\left(|X_{t+h,i} - X_{t,i}|\right), \tag{48}$$

where Q_α is a high historical quantile. A candidate path has transition ratio

$$R_{j,i}(h) = \frac{|\widehat{X}_{j,i}(T+h) - X_{T,i}|}{\max\{E_i(h), f_i\}}, \tag{49}$$

with floor $f_i > 0$. Paths exceeding a hard multiplier are rejected. Remaining paths can receive a soft penalty

$$\Pi_j = \frac{\sum_h \omega_h \frac{1}{2} \sum_{i\in\{Z,V\}} [R_{j,i}(h) - c_{\mathrm{soft}}]_+^2}{\sum_h \omega_h}. \tag{50}$$

This constrains the Zeta–Möbius and volatility trajectories by historically observed state-transition geometry rather than arbitrary clipping.

8 Model-evidence probabilities

Suppose J stable current-conditioned specifications survive all checks. Let S_j denote a composite score combining fit error, trend mismatch, and state-transition extrapolation. Robustly normalize the score by

$$L_j = \left[\frac{S_j - \min_k S_k}{\mathrm{MAD}(S_1, \ldots, S_J)}\right]_+. \tag{51}$$

A weak regularization prior can mildly discourage extreme calibration choices:

$$A_j = \epsilon_M \left|\log \frac{M_j}{\widetilde{M}}\right| + \epsilon_\rho \left|\log \frac{\rho_j}{\widetilde{\rho}}\right|. \tag{52}$$

The model-evidence weight is

$$\boxed{\pi_j = \frac{\exp\left[-(\kappa L_j + A_j)/\tau\right]}{\sum_{k=1}^{J} \exp\left[-(\kappa L_k + A_k)/\tau\right]}.} \tag{53}$$

Proposition 8.1 (Scenario-weight simplex)**.** *The weights in* (53) *satisfy* $\pi_j > 0$ *and* $\sum_j \pi_j = 1$.

Proof. Exponentials are strictly positive and the denominator is the sum of all numerators. □

A concentration diagnostic is the effective number of specifications

$$N_{\mathrm{eff}} = \frac{1}{\sum_{j=1}^{J} \pi_j^2}. \tag{54}$$

It lies in $[1, J]$; equality at J occurs under uniform weights and equality at 1 in the limiting case of complete concentration on one model.

Remark 8.2 (Meaning of "probability"). The π_j are normalized evidence weights over a deterministic model grid. They become empirical predictive probabilities only after a separate calibration exercise demonstrates frequency calibration. They are not risk-neutral probabilities.

9 Trajectory clustering and regime probabilities

Let $\mathcal{H} = \{h_1, \dots, h_m\}$ be selected forecast horizons. For candidate j, define the path-feature vector

$$f_j = \left(\log \frac{P_j(T+h)}{P_T},\, Z_j(T+h) - Z_T,\, \log \frac{\sigma_j(T+h)}{\sigma_T} \right)_{h \in \mathcal{H}}. \tag{55}$$

The feature columns are standardized by weighted means and variances under π_j. A deterministic weighted K-means partitions the candidate set into clusters $C_1, \dots, C_K$.

Definition 9.1 (Regime probability)**.** The model-evidence probability assigned to regime cluster k is

$$p_k = \sum_{j \in C_k} \pi_j. \tag{56}$$

The most likely trajectory cluster is the cluster with largest p_k.

Because the clusters form a partition, $p_k \geq 0$ and $\sum_k p_k = 1$. A medoid trajectory is the member nearest the weighted cluster center and provides a representative path.

The practical regime name combines terminal direction, volatility change, and Zeta–Möbius change. The numerical labels are descriptive summaries rather than distinct stochastic laws.

10 Forecast and risk outputs

10.1 Expected return by horizon

At horizon h, candidate j implies simple return

$$R_j(h) = \frac{P_j(T+h)}{P_T} - 1. \tag{57}$$

The model-evidence mean return is

$$\overline{R}(h) = \sum_j \pi_j R_j(h). \tag{58}$$

The probability of a positive path is

$$p_+(h) = \sum_j \pi_j \mathbf{1}\{R_j(h) > 0\}. \tag{59}$$

10.2 Weighted quantiles and expected shortfall

For $\alpha \in (0,1)$, define the weighted quantile

$$Q_\alpha(h) = \inf \left\{ q : \sum_{j : R_j(h) \leq q} \pi_j \geq \alpha \right\}. \tag{60}$$

The downside expected shortfall is

$$\mathrm{ES}_\alpha(h) = \frac{\sum_{j:R_j(h)\le Q_\alpha(h)} \pi_j R_j(h)}{\sum_{j:R_j(h)\le Q_\alpha(h)} \pi_j}. \tag{61}$$

These provide direct inputs for risk limits.

10.3 Expected volatility and sizing

The all-regime expected annualized volatility is

$$\overline{\sigma}(h) = \sum_j \pi_j \sigma_j(T+h). \tag{62}$$

A simple volatility-targeting multiplier is

$$m(h) = \min\left\{ m_{\max}, \frac{\sigma_*}{\overline{\sigma}(h)} \right\}, \tag{63}$$

where σ_* is a target annualized volatility. This is a research sizing rule, not a universal optimal portfolio solution.

10.4 Zeta–Möbius transition

At a structural horizon h_Z, define

$$\Delta_Z(h_Z) = \overline{Z}(T+h_Z) - Z_T, \tag{64}$$

$$S_Z(h_Z) = \frac{|\Delta_Z(h_Z)|}{s_Z}, \tag{65}$$

$$p_Z^{+}(h_Z) = \sum_j \pi_j \mathbf{1}\{Z_j(T+h_Z) > 0\}, \tag{66}$$

$$p_Z^{\uparrow}(h_Z) = \sum_j \pi_j \mathbf{1}\{Z_j(T+h_Z) > Z_T\}. \tag{67}$$

A zero crossing of the expected Zeta–Möbius path indicates reversal of the signed arithmetic lag balance. It should not be interpreted as a universal price-direction signal without empirical mapping.

11 Trajectory disagreement and forecast confidence

Let $p_1, \ldots, p_K$ be the cluster probabilities. The normalized regime entropy is

$$H_C = -\frac{\sum_{k=1}^K p_k \log p_k}{\log K}, \tag{68}$$

with $H_C \in [0,1]$. Let s_R be the weighted standard deviation of terminal returns and define

$$D_R = 1 - e^{-s_R/c_R}, \tag{69}$$

where c_R is a chosen scale. Directional ambiguity at a short horizon is measured by normalized binary entropy

$$H_2(p) = -\frac{p\log p + (1-p)\log(1-p)}{\log 2}, \tag{70}$$

with $p = p_+(h_0)$. Candidate concentration can be measured by

$$C_N = 1 - \min\left\{1, \frac{N_{\text{eff}}}{N_*}\right\}. \tag{71}$$

The implementation combines these into a disagreement score

$$\mathcal{D} = 100\left(0.40H_C + 0.30D_R + 0.20H_2(p) + 0.10C_N\right), \tag{72}$$

clipped to $[0,100]$. Lower disagreement corresponds to higher forecast confidence.

12 Coherent pathwise state bands

Pointwise quantiles can create visual artifacts. If the subset of paths defining a shortest interval is recomputed independently at each horizon, the lower and upper ribbons may jump even though every underlying trajectory is continuous. To prevent this, the engine selects one fixed subset $\mathcal{C}$ of complete trajectories inside the dominant regime, carrying approximately a chosen mass such as 68%, and uses that same subset at every horizon.

For a state coordinate $Y_j(h)$, define

$$L(h) = \min_{j\in\mathcal{C}} Y_j(h), \qquad U(h) = \max_{j\in\mathcal{C}} Y_j(h). \tag{73}$$

Theorem 12.1 (Continuity of a coherent finite trajectory tube)**.** *If every $Y_j : [0,H] \to \mathbb{R}$ is continuous and $\mathcal{C}$ is a fixed finite index set, then the envelope functions L and U in (73) are continuous.*

Proof. The pointwise minimum and maximum of finitely many continuous real-valued functions are continuous. For two functions, $\max(f,g) = (f+g+|f-g|)/2$ and $\min(f,g) = (f+g-|f-g|)/2$ are continuous. Apply this recursively to finitely many paths. $\square$

Therefore, an unexplained vertical jump in a ribbon built from a fixed set of continuous RDE trajectories is not a credible-band effect. A genuine vertical jump should be generated by an explicit regime-switching or jump mechanism.

13 Historical market-manifold distance

13.1 Feature vector

The engine constructs a local historical feature vector from fifteen components:

normalized log-price state, normalized Zeta–Möbius state, normalized log-volatility state,
5-, 20-, and 60-day momentum,
63- and 252-day drawdown,
5- and 20-day mean return,
20-day realized volatility,
5- and 20-day Zeta–Möbius change,
5- and 20-day log-volatility change.

Features that are not finite at the current date or lack sufficient historical coverage are dropped adaptively, provided a minimum number of usable dimensions remains.

Let $u_t \in \mathbb{R}^d$ denote the retained feature vector. Robustly center and scale historical reference features by componentwise median and MAD:

$$z_t = D^{-1}(u_t - c), \tag{74}$$

where $D = \operatorname{diag}(s_1, \ldots, s_d)$.

Let $\widehat{\Sigma}$ be the covariance of the robustly standardized historical features and define ridge covariance

$$\widehat{\Sigma}_\rho = \widehat{\Sigma} + \rho \bar{d} I, \qquad \bar{d} = \frac{1}{d} \operatorname{tr}(\widehat{\Sigma}), \quad \rho > 0. \tag{75}$$

Proposition 13.1 (Positive-definite manifold metric)**.** *If $\bar{d} > 0$ and $\rho > 0$, then $\widehat{\Sigma}_\rho$ is positive definite even when $\widehat{\Sigma}$ is singular.*

Proof. Because $\widehat{\Sigma}$ is positive semidefinite, for every nonzero v,

$$v^\top \widehat{\Sigma}_\rho v = v^\top \widehat{\Sigma} v + \rho \bar{d} \|v\|_2^2 \geq \rho \bar{d} \|v\|_2^2 > 0.$$

□

Let $\widehat{\Sigma}_\rho = Q \Lambda Q^\top$ and define whitening map

$$W = Q \Lambda^{-1/2}. \tag{76}$$

The whitened Euclidean distance equals a regularized Mahalanobis distance:

$$\|(z - z')W\|_2^2 = (z - z')^\top \widehat{\Sigma}_\rho^{-1} (z - z'). \tag{77}$$

For current point z_T, let $d_{(1)} \le \cdots \le d_{(k)}$ be its k nearest historical distances. Define

$$D_T = \frac{1}{k}\sum_{j=1}^{k} d_{(j)}. \tag{78}$$

Compute the same k-NN distance for every historical reference point. The manifold percentile is

$$\rho_T = \frac{1}{M}\sum_{t=1}^{M} \mathbf{1}\{D_t \le D_T\}. \tag{79}$$

High ρ_T indicates that the current market feature vector lies farther from historical local neighborhoods than most reference observations. This is an anomaly or regime-change warning, not a guarantee of a future price reversal.

14 Explicit regime shifts as a hybrid jump-RDE extension

A visually vertical regime shift should be modeled as a discrete structural event. Let $R_t \in \{1, \dots, K\}$ be a latent regime process and, between regime jumps, let

$$\dot{X}(t) + A_{R_t}\big(X(t) - \mu_{R_t}\big) = G_{R_t}(t). \tag{80}$$

If a transition $i \to j$ occurs at time τ, allow a state reset

$$X(\tau^+) = X(\tau^-) + J_{ij}. \tag{81}$$

The forward regime probabilities may evolve according to

$$\boldsymbol{p}_{h+1} = \boldsymbol{p}_h P_h, \tag{82}$$

where P_h can depend on the current state, the Zeta–Möbius transition, or other evidence. This connects the continuous RDE state to standard regime-switching ideas [10].

Proposition 14.1 (Origin of a vertical state jump)**.** *Suppose the within-regime RDE trajectories are continuous. Then a discontinuity of size J_{ij} at τ occurs if and only if the jump-reset rule (81) is invoked with $J_{ij} \neq 0$. A change in regime probability alone need not make the state path discontinuous.*

Proof. Continuity holds on each open interval between regime transitions by the ODE solution. At τ, (81) gives $X(\tau^+) - X(\tau^-) = J_{ij}$. If $J_{ij} = 0$, the state remains continuous even if the active coefficients change after τ. □

This distinction is operationally important: a vertical plot jump should represent a modeled change in state level, not a change in which trajectories happen to define a pointwise band.

15 What the Zeta–Möbius state is telling us

The Zeta–Möbius state is best interpreted as a *structured memory diagnostic.* The filter asks whether recent observations reinforce or cancel when viewed through a prime-factorization-driven signed lag kernel.

For example,

$$Z_t \approx X_{t-1} - \frac{X_{t-2}}{2^s} - \frac{X_{t-3}}{3^s} - \frac{X_{t-5}}{5^s} + \frac{X_{t-6}}{6^s} + \cdots , \tag{83}$$

while lags $4, 8, 9, 12, \ldots$ vanish because they contain squared prime factors.

Conceptually:

- large $|Z_t|$ means recent data exhibit a strong signed organization under the arithmetic sieve;
- small $|Z_t|$ means stronger cancellation among the selected lag families;
- a rapid change in Z_t indicates reorganization of lag structure;
- a zero crossing indicates reversal of the signed lag balance;
- price movement without a comparable Zeta–Möbius transition can be interpreted as a more local shock, while Zeta–Möbius movement without large contemporaneous price movement can serve as an early structural signal.

The last two statements are hypotheses to be tested empirically, not universal laws. The earlier Zeta–Möbius market-microstructure work found the construction more naturally suited to instability and volatility-state estimation than to direct return prediction [8].

16 Daily RDE Regime Probability Engine

The complete engine can be summarized as the following deterministic workflow.

1. Download and clean adjusted market prices.
2. Form log returns, causal Zeta–Möbius state, and EWMA volatility.
3. Robustly normalize $(\log P, Z, \log \sigma)$.
4. Build a grid of calibration windows, ridge penalties, and recency-decay values; every window terminates at today.
5. For each state coordinate and specification, search $\lambda > 0$ and fit the effective Section 5.2 coefficients by recency-weighted ridge.
6. Apply the exact endpoint-conditioning identity so every trajectory begins at the current observed state.

Table 1: Primary dashboard outputs and their practical interpretation.

Output	Practical use
Most likely trajectory cluster	Current regime classification
Cluster probabilities	Scenario likelihoods / model-evidence mass
Expected return by horizon	Directional forecast
Downside quantiles and expected shortfall	Risk limits
Expected volatility path	Position and option sizing input
Zeta–Möbius transition	Early structural signal
Trajectory disagreement	Forecast confidence
Distance from manifold	Anomaly or regime-change warning

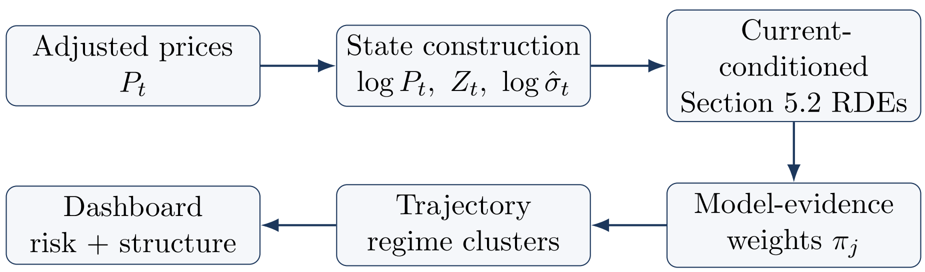

Figure 3: Logical flow of the Zeta–Möbius RDE Regime Probability Engine.

7. Enforce stability, trend-consistency, and historical state-transition-envelope checks.

8. Convert surviving specification scores into model-evidence weights π_j.

9. Cluster complete future trajectories in multi-horizon price/Zeta/volatility feature space.

10. Report regime probabilities, horizon returns, downside quantiles, expected shortfall, volatility paths, Zeta transitions, disagreement, and historical-manifold distance.

17 Interpretive and statistical limitations

Several distinctions are essential.

Model-evidence probability is not calibrated probability. The scenario weights in (53) rank deterministic specifications by relative fit and plausibility. Calibration against realized future regime frequencies is required before interpreting p_k as a frequentist probability statement.

Geometric central price path is not the full conditional price mean. The RDE mean is fitted to log price. The displayed central path is $\exp(\mathbb{E}[L])$. Since exp is convex,

$$\exp(\mathbb{E}[L]) \leq \mathbb{E}[e^{L}], \tag{84}$$

with strict inequality when L has nonzero conditional variance. The engine therefore supplies a deterministic geometric central scenario, not a full lognormal pricing distribution.

The Zeta–Möbius sign is not intrinsically directional. The sign encodes which signed lag family dominates. Any bullish/bearish interpretation must be learned from historical conditional outcomes for the chosen signal and parameters.

Polynomial Wiener forcing is generally nonstationary. The Section 5.2 expectation itself shows that even a stable linear kernel does not guarantee a stationary mean when the forcing is a polynomial of a Wiener level. This is why finite-horizon conditioning, low polynomial degree, and transition-envelope control are important.

Regime shifts require an explicit mechanism. A jump in a visualization is meaningful only if it follows from a jump/reset model, a discrete observation update, or a true discontinuity. A pointwise credible-band algorithm should not be allowed to create artificial structural breaks.

18 Conclusion

The Zeta–Möbius RDE Regime Probability Engine separates market-state construction from market-state evolution. The Möbius transform supplies a deterministic arithmetic sieve over recent history. The structured-forcing RDE supplies explicit, current-conditioned trajectories without Monte Carlo sampling. A grid of calibration specifications produces deterministic model uncertainty, which is converted into model-evidence weights and clustered into scenario regimes. The resulting outputs map directly to practical decisions: current regime classification, scenario likelihoods, directional forecasts, downside risk limits, volatility-based sizing, structural Zeta–Möbius transitions, disagreement-based confidence, and historical-manifold anomaly warnings.

The framework's most important mathematical safeguards are also conceptual safeguards: every forecast is anchored exactly at the current observation; strongly extrapolative state paths are penalized against historical transition geometry; coherent uncertainty tubes use a fixed set of complete paths; and genuine regime jumps are reserved for explicit hybrid transition mechanisms. These distinctions keep the dashboard's visual outputs aligned with the mathematics that generates them.

A Expanded derivation of the Wiener kernel

Starting from

$$J_q(t;\lambda) = e^{-\lambda t}\int_0^t e^{\lambda u}u^q\,\mathrm{d}u,$$

integration by parts gives

$$\int_0^t e^{\lambda u} u^q \, \mathrm{d}u = \frac{e^{\lambda t} t^q}{\lambda} - \frac{q}{\lambda} \int_0^t e^{\lambda u} u^{q-1} \, \mathrm{d}u.$$

Multiplying by $e^{-\lambda t}$ yields the recurrence (23). Repeating q times produces

$$\begin{aligned} J_q(t;\lambda) = \frac{t^q}{\lambda} - \frac{q t^{q-1}}{\lambda^2} + \frac{q(q-1)t^{q-2}}{\lambda^3} - \cdots \\ + (-1)^q \frac{q!}{\lambda^{q+1}} + (-1)^{q+1} \frac{q!}{\lambda^{q+1}} e^{-\lambda t}, \end{aligned}$$

which is exactly (24).

Differentiating the integral definition also gives the useful ODE identity

$$\frac{\partial}{\partial t} J_q(t;\lambda) + \lambda J_q(t;\lambda) = t^q, \qquad J_q(0;\lambda) = 0. \tag{85}$$

This identity is what leads directly to the conditioned derivative formula (32).

B Why current conditioning preserves the forcing law

Write the unconditioned mean as $m = m_p + m_h$, where m_p is any particular solution of (29) and $m_h(t) = Ce^{-\lambda t}$ is the homogeneous solution. Conditioning at T merely changes the homogeneous coefficient:

$$\begin{aligned} m^{(T)}(t) &= m(t) + e^{-\lambda(t-T)}[x_T - m(T)] \\ &= m_p(t) + e^{-\lambda t}\Big(C + e^{\lambda T}[x_T - m(T)]\Big). \end{aligned}$$

Thus the polynomial forcing $g(t)$ is not altered. Only the homogeneous initial-condition component is updated. This is why the correction is exact rather than an ad hoc vertical translation of the whole forecast.

C Continuity versus a modeled regime jump

Suppose $\{Y_j\}_{j\in\mathcal{C}}$ is the fixed coherent set used for the scenario ribbon. By Theorem 10.1, $L(h)$ and $U(h)$ are continuous if all Y_j are continuous. If the hybrid model introduces a jump J_{ij} at τ for every trajectory in a subset $\mathcal{C}_{ij}$, then for those trajectories

$$Y_j(\tau^+) - Y_j(\tau^-) = J_{ij,Y}.$$

The envelope may therefore jump because the underlying trajectories themselves jump. This is the correct mechanism for a visual regime-reset discontinuity.

D Default research configuration

The current research implementation uses values in the neighborhood of those shown in table 2. They are engineering defaults, not universal constants.

Table 2: Representative default settings in the current engine.

Component	Parameter	Default
Zeta–Möbius	lag depth N	100
Zeta–Möbius	exponent s	1.5
Forecast	horizon	100 trading days
RDE	polynomial degree	2
Calibration	window range	126–756 trading days
Calibration	λ grid	0.15–35 annual units
Regimes	number of clusters	5
Risk	downside probability	10%
State bands	coherent target mass	68%
Transition envelope	historical quantile	97.5%
Manifold	reference observations	600
Manifold	nearest neighbors	10
Manifold	covariance ridge	0.10

References

[1] T. M. Apostol, *Introduction to Analytic Number Theory*, Springer, 1976.

[2] H. Davenport, *Multiplicative Number Theory*, 3rd ed., Springer, 2000.

[3] P. P. Vaidyanathan, "Srinivasa Ramanujan and signal-processing problems," *Philosophical Transactions of the Royal Society A*, 2020.

[4] T. T. Soong, *Random Differential Equations in Science and Engineering*, Academic Press, 1973.

[5] L. Arnold, *Random Dynamical Systems*, Springer, 1998.

[6] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications*, Springer, 2003.

[7] J. Riposo and P. Zavari, *Random Differential Equations with Structured Stochastic Forcing*, SSRN Working Paper No. 6193198, 2026. https://doi.org/10.2139/ssrn.6193198.

[8] P. Zavari, *Zeta–Möbius Market Microstructures: A Finance Extension of the Zeta Random-Walk Construction with Benchmarks*, working paper, 2026.

[9] F. Black and M. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81(3), 637–654, 1973.

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

[11] J. MacQueen, "Some Methods for Classification and Analysis of Multivariate Observations," in *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 1967.

[12] P. J. Huber and E. M. Ronchetti, *Robust Statistics*, 2nd ed., Wiley, 2009.