

Cylindrical Yield-Curve Dynamics and the Arbitrage-Free Class*

Zeyu Cao[†]

Shaosai Huang[‡]

September 15, 2026

Abstract

Finite-factor term-structure models assign no independent risk to exposures between their factors, although desks treat them as risks of their own. The cylindrical term-structure model of [13] gives each maturity an independent shock, but specifies real-world dynamics, leaving the link to arbitrage pricing open. This paper settles that question for every finite truncation of the deformation at fixed spread intensity, and then restores the stochastic spread in two conventions. The key observation is that the deformation’s damped sine modes are imaginary parts of complex exponentials, so the smallest shift-invariant family containing them is a doubled sine–cosine envelope whose added coordinates carry no noise. On this envelope, whether the factor drift is arbitrage-free reduces to one polynomial divisibility, which the model’s maturity diffusion fails on every mode. Because the failure sits in the noiseless coordinates, no change of measure removes it. No such truncation admits an equivalent martingale measure; for generic maturities an explicit arbitrage uses one bond more than there are modes; and the defect grows with resolution. Giving those coordinates noise of their own turns the impossibility into a market price of risk that diverges with resolution. Adjoining those coordinates, keeping the model’s noise and correcting only the drift makes every truncation arbitrage-free. The required convexity has an exact limit under refinement, with a short-end jump resolved in a boundary layer at the resolution scale. Bond-level volatilities converge, but forward-curve distributions do not: within this construction, maturity-local noise, no-arbitrage at every resolution and a curve-valued limit cannot coexist. The short-end value survives the stochastic spread. It is conserved at every vol-of-vol in the uncorrelated affine convention, and recovered at zero vol-of-vol in the bilinear quadratic–Gaussian convention under a recovery normalization, where each truncation has an exact existence criterion. The consistency test and the arbitrage construction extend beyond this model.

Keywords. Term structure of interest rates; Heath–Jarrow–Morton drift condition; finite-dimensional realisations; quasi-exponential curves; complex spectrum; maturity-white risk; bond arbitrage; quadratic–Gaussian models.

MSC 2020. 91G30, 60H15, 91G80. **JEL.** C65, E43, G12, G13.

*Working paper. Comments welcome.

[†]Independent Researcher, Long Island City, NY 11101, USA. Email: josephcao891011@gmail.com.

[‡]Kspectra Research Inc., Toronto, ON, M2N 0G3, Canada. Email: arthur.foxie.huang@kspectra.ai.

1 Introduction

A finite-factor term-structure model assigns no independent risk to exposures that fall between its factors. For a symmetric butterfly, the modelled exposure is only the second difference of smooth factor loadings, so it vanishes quadratically as the wings are brought in towards the centre maturity. Desks nevertheless hedge and charge for the position as though it carried local risk of its own. In practice, this gap is managed rather than eliminated. Factor truncation, tenor buckets, proxy hedges, convexity overlays and basis-risk reserves make a finite production system workable. They are operational accommodations: they do not specify how the residual maturity risk should enter a common dynamic pricing model.

Turning that residual-risk problem into a production pricing model requires answers to two separate questions. *Finite implementation* asks what a finite truncation preserves and what depends on its maturity resolution. *Pricing consistency* asks whether deterministic maturity roll-down and the Heath–Jarrow–Morton (HJM) convexity drift remain inside the volatility span, so that the truncation admits an equivalent local martingale measure. The first question concerns representation; the second concerns no-arbitrage.

The cylindrical method [13] addresses the risk representation first: one innovation is assigned to each maturity, rather than imposing smooth interpolation across neighbouring tenors. Write $\theta \in [0, \theta^*]$ for time to maturity up to the maximum modelled maturity θ^* . The forward curve is carried by a deformation field X ,

$$f_t(\theta) = r_t + s_t[Y(\theta) + X_t(\theta)],$$

with r_t a short rate, s_t a scalar stochastic spread and Y a deterministic shape, and X solves the second-order stochastic evolution equation

$$dX_t = \left[\partial_\theta + \frac{\kappa}{2} \partial_\theta^2 \right] X_t dt + \sigma_0 d\mathcal{B}_t, \quad X_t(0) = X_t(\theta^*) = 0.$$

Here ∂_θ transports the curve as calendar time passes, $\frac{\kappa}{2} \partial_\theta^2$ diffuses shocks along the maturity axis at scale κ . The noise $\mathcal{B}$ is a *cylindrical* Brownian motion on $H = L^2([0, \theta^*], e^{2\theta/\kappa} d\theta)$ of intensity σ_0 : one independent innovation per maturity, which is the whole point of the construction. The Dirichlet realisation of that generator has orthonormal eigenbasis

$$e_n(\theta) = \sqrt{2/\theta^*} e^{-\theta/\kappa} \sin \omega_n \theta, \quad \omega_n = \frac{n\pi}{\theta^*}, \quad \lambda_n = \frac{1 + \kappa^2 \omega_n^2}{2\kappa},$$

with eigenvalues $-\lambda_n$, so the coordinates $x_n = \langle X_t, e_n \rangle_H$ are independent Ornstein–Uhlenbeck processes reverting at rate λ_n . Each e_n is a damped sine of maturity frequency ω_n and common damping scale κ . Equivalently, it is the imaginary part of a complex exponential with exponent $-1/\kappa + i\omega_n$: the real sine–cosine pair carries the conjugate pair, and forward rates and bond prices remain real. These shapes belong to the class associated with finite-dimensional arbitrage-free realisations [6, 19]. The dynamics, however, are specified under the real-world probability measure $\mathbb{P}$, while the link to the risk-neutral pricing measure $\mathbb{Q}$ is left open. With the spread held fixed, the link cannot be supplied as written: no finite sine truncation of the deformation admits an equivalent local martingale measure with its own volatility. That is the problem solved first; §6 then restores the stochastic spread.

We call X the *cylindrical deformation field* and the displayed specification the *source model*. Unqualified, “the model” means the source model. In §§3–5, where the spread is held fixed, it

means that model's deformation sector, truncated to its first N sine modes. “The completion” means the arbitrage-free construction of §5; each stochastic-spread family is named by the assumption and convention that defines it.

What makes the question answerable. To carry the analysis past the point at which the link to $\mathbb{Q}$ was left open we start from the standard route: consistency is decided on finite-dimensional shift-invariant families of curve shapes [6, 7, 19], in the invariant-manifold form of [21]. Here that route is compulsory rather than convenient. Under cylindrical noise the forward rate at a fixed maturity is not a semimartingale (§3), so the drift condition cannot be posed maturity by maturity at all. On the curve the comparison is not merely hard but ill-posed; finite-dimensional invariant subspaces are where the question can be made rigorous.

The route is blocked at its first step. It asks for a shift-invariant family, and the model has none. Its N sine modes span no invariant space, so there is nothing for the realisation machinery to act on — which is where the matter had rested.

The block dissolves once the spectrum is read as complex. Each mode $e^{-\theta/\kappa} \sin \omega_n \theta$ is the imaginary part of $e^{(-1/\kappa + i\omega_n)\theta}$. A conjugate pair admits no one-dimensional real invariant subspace, so the smallest invariant space containing one mode is two-dimensional. The model's span therefore completes to the $2N$ -dimensional sine-cosine *envelope*, and it is on that envelope — not on the model's own state space — that the comparison can be made. This is the step the rest of the paper rests on, and it pays twice over. It supplies the invariant family the standard route was asking for. And the coordinates it adjoins carry no noise, while the missing rotation ω_n acts precisely on them: what would have looked, on the model's own span, like a correctable disagreement about drift is revealed as an obstruction that no change of measure can remove. Complexification is the mechanism here, not a change of notation.

The remaining steps follow from the envelope. On it the deformation and the arbitrage-free transport both act by matrices, so the comparison of drifts becomes algebra. A generator with symbol p induces the factor drift $p(M)^\top$ where no-arbitrage requires $M^\top$, and the two agree exactly when $m_M \mid p - z$: the comparison is settled by one polynomial divisibility, and the convexity that HJM adds is carried by the deterministic sector. The envelope is also an ordinary finite-dimensional family, so the obstruction can be repaired rather than merely reported. Keeping the model's loadings and its noise and correcting only the drift turns every truncation into an arbitrage-free model, and the refinement limit of those models fixes the resulting convexity exactly.

What this paper establishes. The deformation generator $\partial_\theta + \frac{\kappa}{2} \partial_\theta^2$ has symbol $p(z) = z + \frac{\kappa}{2} z^2$, so the symbol test fails on every decaying or oscillatory block; inside the weighted state space the tolerated class is the constants alone (§4). The doubling has a second consequence. The adjoined cosine partners violate the Dirichlet conditions, so the envelope leaves the domain of the closed generator (Proposition 4.1); the comparison is therefore a tangency statement about the formal symbol, which is precisely why no formulation inside the original domain is available. For the sine truncations at fixed intensity the failure is exact and measure-invariant at every order. No such truncation admits an equivalent local martingale measure (Remark 4.4); the witness construction exhibits the failure as a bond arbitrage (Theorem 4.6); and the obstruction modulus shows the unreachable part of the drift to be extensive rather than washed out by refinement (Proposition 4.8). For the infinite model the repairing shift fails the Girsanov finite-energy condition (Proposition 4.10), and the fully infinite bond-market statement is posed as open. When the silent coordinates are given noise of their own (§4.7), the impossibility

becomes a price: every truncation then admits an equivalent measure, at a Sharpe cost growing like $N^{3/2}$, and the infinite factor laws are mutually singular (Proposition 4.11). Pointwise rates are not semimartingales (§3), so the finite-dimensional route is the rigorous one. The minimal completion has exact first-span interior limit $q^\infty = \frac{\sigma_0^2}{2} e^{-2\tau/\kappa}$, a boundary layer on the scale θ^*/N with first-layer profile peak $1.455 \times \sigma_0^2/2$, and a short-end jump $\sigma_0^2/2$ that is the drift cost of that completion. No curve-valued limit exists (Proposition 5.4). In the uncorrelated case, the cost survives the CIR-intensity extension at every vol-of-vol and the quadratic–Gaussian one on its $\eta = 0$ slice. For CIR intensity, the infinite-span surrogate exists precisely when $\lambda \geq \eta\sigma_0\sqrt{\kappa/2}$ (§6.1); and the quadratic–Gaussian bilinear channel, under a stated recovery normalization, has an exact finite- N existence criterion (Proposition 6.16).

In the table below, N is the number of retained maturity modes and θ^*/N is their maturity-resolution scale. For each fixed N , the original truncation and its completion are ordinary finite-factor bond markets; refinement in N measures sensitivity to maturity resolution.

Issue	Common operational patch	Maturity-local method	Resolution in this paper
Finite implementation	Variance-ranked factors, tenor buckets and local convexity overlays. These make a finite system usable, but leave resolution as an economically meaningful modelling choice.	One innovation per maturity restores local risk before truncation, but the cylindrical state itself is infinite-rank and cannot be evolved exactly by a production system.	For every N , an ordinary finite HJM completion; an exact first-span convexity limit and θ^*/N boundary layer; and non-tight curve laws under refinement (Theorem 5.2, Proposition 5.4).
Pricing consistency	Proxy hedges in nearby bonds, bucketed greeks, and basis-risk reserves and limits. They manage residual P&L, but do not themselves impose a joint martingale measure.	The real sine modes capture maturity-local shocks under $\mathbb{P}$, while deterministic roll-down generates missing cosine directions outside the original shock span.	An exact range obstruction; a generic $(N + 1)$ -bond arbitrage witness; an extensive stationary defect; and a completion that adjoins the roll-down states and required HJM convexity (Theorem 4.3, Theorem 4.6, Proposition 4.8).

These columns answer different questions. Operational patches manage portfolios in a finite system; the maturity-local method supplies a richer statistical risk field; the completion determines whether each finite truncation can price bonds without arbitrage and quantifies the cost of doing so.

Once a volatility and the minimal envelope have been selected, HJM consistency determines the drift uniquely within that completion. That drift is the cost quantified below; the premium a market attaches to the original unspanned risk is a separate question, which no-arbitrage alone does not settle.

Methodological contributions: four tools. The paper is organised so that four of its results can be lifted out of the instance that motivates them.

(i) *A symbol test for consistency* (Theorem 4.3(a)). The criterion is not tied to this instance. For any constant-coefficient generator with symbol p and any shift-invariant block with generating matrix M , the arbitrage-free drift is carried if and only if $m_M \mid p - z$. For diagonalisable M this says that the symbol must fix the block’s spectrum pointwise, which is the form in which it is usually easiest to apply.

(ii) *A constructive arbitrage witness* (Theorem 4.6, Remark 4.7). From a state-linear drift discrepancy lying outside the volatility range, an explicit self-financing admissible arbitrage in $N + 1$ fixed-maturity bonds, on an open dense set of maturity tuples. The construction needs only analytic loadings, a Girsanov-unreachable discrepancy, and a nondegenerate state — and it exhibits the portfolio, where the fundamental theorem of asset pricing gives only existence.

(iii) *A quantitative obstruction modulus* (Proposition 4.8). The quotient distance from the discrepancy to the volatility range, with its exact stationary law: not whether a model violates no-arbitrage but *by how much*, and whether refinement attenuates the violation. Asking that question is not new: it is what the Hansen–Jagannathan distance asks in stochastic-discount-factor space [24]. This modulus is a narrower, model-internal object — a quotient distance in declared factor coordinates whose stationary law is exact. In that normalization it is a diagnostic, not a price.

(iv) *Minimal completion as a repair, with a recoverable cost* (§5, §§6.1–6.5). The repair above, stated as a method rather than a computation: faced with a model that fails no-arbitrage, complete its state to the minimal shift-invariant hull, correct only the drift, and read the cost off the refinement limit instead of recording the failure and stopping. Repairing rather than discarding an inconsistent curve family is the move of the affine arbitrage-free Nelson–Siegel class [9], which keeps the Nelson–Siegel loadings and restores consistency through an explicit yield-adjustment term. What differs here is that the family is forced rather than given — the minimal hull of a state space possessing none — and that the cost is read off a refinement limit rather than at a fixed factor count. With the sine-only noise kept, the same short-end coefficient reappears in the uncorrelated affine convention at every vol-of-vol and in the bilinear convention on its zero-vol-of-vol slice. Across completions the cost can move: isotropic loading doubles the jump and smoothed cutoffs alter the overshoot (Remark 5.3(iii)). The positive-vol-of-vol bilinear case is not established.

Related work. The drift restriction is [25] in the parametrisation of [36]. Which curve families admit finite-dimensional arbitrage-free realisations was settled by Björk–Christensen, Björk–Svensson and Filipović [6, 7, 19, 20], with the invariant-manifold formulation in [21]. That classification is indifferent to whether the generating spectrum is real, and this paper takes up the half of it that term-structure practice has left unused. On a complex block the invariant family is two-dimensional, so a model built from sine modes alone is not merely inconvenient but incomplete; the coordinate its completion adjoins is where the obstruction turns out to live. Where the classification gives a condition to be verified family by family, Theorem 4.3(a) gives a decision procedure for the factor drift: one divisibility of polynomials.

Random-field term structures with *smooth* spatial covariance [33, 22, 39] carry a drift condition written in terms of the covariance kernel. Cylindrical noise is the singular limit of that family: the kernel degenerates to a delta and the kernel-level drift condition is no longer available. Theorem 5.2 supplies what replaces it, and Remark 5.3(i) shows the resulting convexity to be exactly the half-delta that heuristic predicts — now with a boundary layer, a rate, and an exact constant. The analytic background for infinite-dimensional formulations is [8, 15], and the pointwise regularity of §3 transcribes stochastic-heat-equation results [41, 40]. The use made of them here is one those results are not usually asked to support: the pointwise drift condition is unavailable, so the finite-dimensional route is forced rather than chosen. Against the long-end constraint [16, 27], Proposition 5.5 separates the two effects: complexification neither causes nor cures the classical pathology.

The equivalence–singularity dichotomy of §4.7 rests on a classical engine: Kakutani’s theorem

for product measures [30] and, for Gaussian evolution equations, the criteria of Huebner and Rozovskii [28] (surveyed in [11]), under which the laws generated by two drift operators are equivalent or mutually singular according to the order of the operators' difference. What is new is the financial reading. The invariant that decides the dichotomy is a market-price-of-risk energy, and the model sits on the singular side by a margin that grows like the cube of the truncation order. This bears on a question raised by [3], who study parabolic perturbations of the Musiela equation with the model in view and ask for the amount of arbitrage such a perturbation introduces to be quantified. Their metric is the pricing error of the perturbed curve; the defect energy of §4.7 is a complementary one, and on the envelope it answers as follows: every nonzero parabolic coefficient makes the infinite-envelope energy infinite, while every fixed finite truncation remains equivalent after tilting, at a Sharpe cost that grows with the truncation order. The heat-modulated model of [32] moves the Laplacian out of the drift and into the covariance process precisely in order to stay inside the HJM class; Proposition 4.11(c) shows that a nonzero second-order term cannot remain in the factor drift while preserving equivalence to the arbitrage-free infinite-envelope factor law.

The affine and quadratic–Gaussian classes of §§6.1–6.5 are [18, 14] and [2, 34, 10], where solvability of the loading Riccati over the maturity range of interest is typically assumed. Here it is decided: for the affine surrogate by the exponential-moment criterion of [4] (Theorem 6.5), and for the quadratic–Gaussian channel at every fixed truncation by a conjugate-point criterion (Proposition 6.16), whose consequences uniform in N are developed in the companion manuscript [42]. That finitely many factors need not span fixed-income risk is an established empirical theme [12, 26], and [17] bounds which spectral statements are testable under $\mathbb{P}$. The framework of [35] is the natural one in which to relax deterministic factor volatility, and Markovian lifts of rough volatility [1] generate the same quasi-exponential structure on a different curve.

Set against that background, what this paper contributes begins at a single observation: the coordinate the doubled envelope adjoins is silent, and the obstruction lives precisely there. The decision procedure, the bond witness and the modulus are what that observation makes usable; the exact completion limit with its boundary layer (Theorem 5.2), and the failure of the repaired family to converge as a curve model (Proposition 5.4), are what it costs. We have not found the finite-portfolio construction elsewhere.

Organization. §2 assembles the yardstick: the state space and curve dynamics, the drift restriction with proof, the quasi-exponential classification, and the consistent affine realisations. §3 states the model and closes the pointwise route: under cylindrical noise, forward rates are not semimartingales, so the comparison is conducted on finite-dimensional invariant subspaces. §4 conducts it at three levels — the symbol test in factor coordinates, the bond-market witness, and the obstruction modulus — asks what survives in the infinite model, and prices the obstruction once the silent coordinates are loaded. §5 builds the completion, computes its singular limit and its drift cost, shows that no curve-valued limit exists, and disposes of the long end. §6 restores the stochastic spread in its two conventions, affine and bilinear, and shows that the short-end cost survives it — at every vol-of-vol in the affine convention, on the zero-vol-of-vol slice in the bilinear one. §7 concludes: it answers the two questions above, gives the financial reading of the answers, and states what remains open.

Companion material. The boundary theory of the quadratic–Gaussian existence region is developed separately in [42]. The fixed- N criterion and the principal financial conclusions proved here are self-contained; the uniform-in- N positive-vol-of-vol region quoted after Proposition 6.16

is a result of the companion manuscript.

Scope relative to the source model. The source construction models real-world dynamics, states that the arbitrage link is future work, and truncates the maturity domain at finite $\theta_{\max}$ explicitly to avoid the $\theta \rightarrow \infty$ pathologies of Dybvig–Ingersoll–Ross type [13]. Nothing here identifies an error in that statistical construction; the question is what it costs to turn the method into a model under $\mathbb{Q}$, and the answer has exact structure.

AI-use disclosure. The authors used Anthropic Claude Code and OpenAI Codex as interactive research and writing tools. They assisted with exploratory discussion, testing and refinement of ideas, literature and source organization, code development and verification, mathematical error checking, and editorial revision.

2 The arbitrage-free class

The introduction posed the question and wrote down the model; before the model can be measured, the standard has to be fixed. This section assembles that yardstick. Four ingredients are fixed once and used everywhere: the state space and the curve dynamics; the no-arbitrage drift restriction, with proof; the classification of curve shapes compatible with finite-dimensional arbitrage-free realisations; and the characterisation of the consistent realisations themselves. Nothing here is specific to the cylindrical model. A reader who has the HJM drift condition and the quasi-exponential class at hand may skim to §3.

2.1 State space and curve dynamics

Everything is set on one Hilbert space of curves, chosen so that maturity roll-down is a strongly continuous semigroup and evaluation at a maturity is continuous.

State space and shift semigroup. Fix a weight $w : \mathbb{R}_+ \rightarrow [1, \infty)$, nondecreasing, with $\int_0^\infty w^{-1/3} < \infty$, and let H_w be the Filipović space [19] of absolutely continuous $h : \mathbb{R}_+ \rightarrow \mathbb{R}$ with

$$\|h\|_w^2 := |h(0)|^2 + \int_0^\infty |h'(\tau)|^2 w(\tau) d\tau < \infty.$$

H_w is a separable Hilbert space of continuous curves; point evaluations are continuous; the shift semigroup $(S_t h)(\tau) := h(\tau + t)$ is strongly continuous on H_w with generator ∂_τ on its natural domain. Membership of a quasi-exponential is an additional integrability requirement, not a consequence of $\operatorname{Re} \lambda \leq 0$: a term $p(\tau)e^{\lambda\tau}$ belongs to H_w precisely when its derivative is square-integrable against w . In particular constants belong to H_w , whereas undamped oscillations and nonconstant polynomials do not (because $w \geq 1$). Finitely many strictly damped modes belong to H_w whenever w is chosen to grow slowly enough relative to their damping; this includes the model’s modes, for example with $w(\tau) = e^{a\tau}$ and $0 < a < 2/\kappa$. The model’s curves live on $[0, \theta^*]$; its eigenmodes extend uniquely to $\mathbb{R}_+$ by their quasi-exponential expressions, and all comparisons below are made on these extensions. The role of the Dirichlet boundary conditions is treated separately in Proposition 4.1.

Whenever an envelope or completion is read as an H_w -valued curve model, we also impose the *damping-compatible* condition

$$C_w := \int_0^\infty e^{-2\tau/\kappa} w(\tau) d\tau < \infty, \tag{1}$$

which the exponential weight above satisfies. The condition is used only where loading blocks are embedded as curves in H_w ; statements made in factor coordinates, and the product-law statements of §4.7, stand without it.

HJMM dynamics. Under a measure $\mathbb{Q}$ carrying a d -dimensional Brownian motion $\mathbf{W}$, the forward curve $r_t \in H_w$ follows the Heath–Jarrow–Morton–Musiela dynamics

$$dr_t = (\partial_\tau r_t + \alpha_t) dt + \sigma_t d\mathbf{W}_t, \quad (2)$$

interpreted in the mild sense $r_t = S_t r_0 + \int_0^t S_{t-s} \alpha_s ds + \int_0^t S_{t-s} \sigma_s d\mathbf{W}_s$. Here $\alpha : \Omega \times \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}$ and σ is row-vector valued; both are jointly measurable, adapted and locally integrable in the appropriate norms. (See [19, 15] for the existence theory; none of the delicate cases arise for the affine models used here, whose solutions are explicit.)

2.2 The drift condition, with proof

The restriction below is the Heath–Jarrow–Morton no-arbitrage condition [25] in the Musiela parametrisation [36]; we record it with proof because every later statement is an instance of it.

Lemma 2.1 (HJM drift restriction). *Let bond prices be $P(t, T) = \exp(-\int_0^{T-t} r_t(u) du)$ and let the savings account be $B_t = \exp(\int_0^t r_s(0) ds)$. Suppose r solves (2) and that, on every bounded maturity triangle, the integrability conditions for stochastic Fubini and Itô’s formula hold. Assume also that, for $\mathbb{Q} \otimes dt$ -a.e. (ω, t) , the maturity maps $\tau \mapsto \alpha_t(\tau)$ and $\tau \mapsto \sigma_t(\tau)$ are locally integrable, so their maturity integrals are absolutely continuous. Then $P(\cdot, T)/B$ is a local $\mathbb{Q}$ -martingale for every T if and only if*

$$\alpha_t(\tau) = \sigma_t(\tau) \int_0^\tau \sigma_t(u)^\top du \quad \text{for a.e. } \tau, \mathbb{Q} \otimes dt\text{-a.e.} \quad (3)$$

Under the physical measure $\mathbb{P}$ the same computation gives $\alpha_t(\tau) = \sigma_t(\tau) (\int_0^\tau \sigma_t(u)^\top du + \gamma_t)$ with γ the market price of risk, which is the form quoted in [13], Eq. (2); a proof is included for completeness.

Proof. Work in the HJM parametrisation $f(t, T) = r_t(T - t)$, so that $df(t, T) = \alpha^H(t, T) dt + \sigma^H(t, T) d\mathbf{W}_t$ with $\alpha^H(t, T) = \alpha_t(T - t)$, $\sigma^H(t, T) = \sigma_t(T - t)$; the ∂_τ term in (2) is absorbed by the reparametrisation. Let $Y_t := \log P(t, T) = -\int_t^T f(t, u) du$. By stochastic Fubini,

$$dY_t = f(t, t) dt - \int_t^T df(t, u) du = \left(r_t(0) - A(t, T) \right) dt - \Sigma(t, T) d\mathbf{W}_t,$$

with $A(t, T) = \int_t^T \alpha^H(t, u) du$ and $\Sigma(t, T) = \int_t^T \sigma^H(t, u) du$. Itô on $P = e^Y$:

$$\frac{dP(t, T)}{P(t, T)} = \left(r_t(0) - A(t, T) + \frac{1}{2} |\Sigma(t, T)|^2 \right) dt - \Sigma(t, T) d\mathbf{W}_t.$$

The discounted bond is a local martingale iff the drift equals $r_t(0)$, i.e. $A(t, T) = \frac{1}{2} |\Sigma(t, T)|^2$ for all T . Differentiating in T gives $\alpha^H(t, T) = \sigma^H(t, T) \Sigma(t, T)^\top$, which is (3) in Musiela coordinates $\tau = T - t$. Reversing the steps gives the converse. Under $\mathbb{P}$, replace “local martingale” by “local martingale after the Girsanov tilt $d\mathbf{W}^\mathbb{Q} = d\mathbf{W}^\mathbb{P} + \gamma_t dt$.” $\square$

The condition is on the *drift given the volatility*; the volatility itself — and in particular its spectrum in the sense of §2.3 — is unconstrained. The rest of the paper exploits this division of labour.

2.3 Quasi-exponential structure

In Musiela coordinates, calendar time rolls every factor loading down the maturity axis. A finite-state term-structure model is closed only when that roll-down remains inside the chosen loading span. This is the financial meaning of shift invariance.

Definition 2.2. $\text{QE} := \{\sum_i p_i(\tau) e^{\lambda_i \tau} : p_i \in \mathbb{C}[\tau], \lambda_i \in \mathbb{C}\}$, real elements understood as those equal to their conjugates.

Theorem 2.3 (Finite-dimensional shift-invariant subspaces). *Let $V \subset C^1(\mathbb{R}_+, \mathbb{R})$ be finite-dimensional with $\partial_\tau V \subseteq V$ (elements are then automatically C^∞). Let M be the matrix of $\partial_\tau|_V$ in a basis $\mathbf{B} = (b_1, \dots, b_n)^\top$, i.e. $\mathbf{B}' = M\mathbf{B}$. Then:*

- (a) $\mathbf{B}(\tau) = e^{M\tau} \mathbf{B}(0)$, and every $v \in V$ satisfies the constant-coefficient ODE $\chi_M(\partial_\tau)v = 0$, where χ_M is the characteristic polynomial of M ; hence $V \subseteq \text{QE}$ with exponents $\text{spec } M$ and polynomial degrees bounded by the sizes of the Jordan blocks. Conversely a finite-dimensional quasi-exponential space is shift-invariant when, for each exponent λ , it contains complete downward-closed Jordan chains $\{e^{\lambda\tau}, \tau e^{\lambda\tau}, \dots, \tau^{m_\lambda-1} e^{\lambda\tau}\}$ (and the conjugate chains in the real case). An arbitrary subspace of QE need not be shift-invariant.
- (b) (Real form.) $\mathbb{R}^n$ decomposes into M -invariant blocks on which M is either a real Jordan block or, for a nonreal pair $\mu \pm i\omega$, a real Jordan block with diagonal 2×2 cells $C = \begin{pmatrix} \mu & \omega \\ -\omega & \mu \end{pmatrix}$ and identity super-cells. On a semisimple complex cell,

$$e^{C\tau} = e^{\mu\tau} R(\omega\tau), \quad R(\phi) = \begin{pmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{pmatrix}, \quad (4)$$

so the basis pair is $e^{\mu\tau}(\cos \omega\tau, \sin \omega\tau)$: exponential decay times rotation — precession.

- (c) (Integration.) $\mathbf{I}(\tau) := \int_0^\tau \mathbf{B}(u) du = (\int_0^\tau e^{Mu} du) \mathbf{B}(0)$, where on the invertible part of M , $\int_0^\tau e^{Mu} du = M^{-1}(e^{M\tau} - \mathbf{1})$, and on the nilpotent part ($M = N$, $N^m = 0$), $\int_0^\tau e^{Nu} du = \sum_{k=0}^{m-1} \frac{\tau^{k+1}}{(k+1)!} N^k$. In particular QE is closed under $\int_0^\tau$, at the price of raising polynomial degree on the kernel of M .
- (d) (Products.) $p_1 e^{\lambda_1 \tau} \cdot p_2 e^{\lambda_2 \tau} = p_1 p_2 e^{(\lambda_1 + \lambda_2) \tau}$, so QE is closed under multiplication, with exponent set closed under addition.

Part (a) is the algebraic core of the Björk–Christensen consistency classification [6, 19], and (d) underlies the exponent-doubling phenomenon in strict finite-dimensional realisations [7].

Proof. (a) $\mathbf{B}' = M\mathbf{B}$ is a linear constant-coefficient system, so $\mathbf{B}(\tau) = e^{M\tau} \mathbf{B}(0)$; entries of $e^{M\tau}$ are quasi-exponentials by the Jordan form. By Cayley–Hamilton, $\chi_M(\partial_\tau)$ annihilates every component of $\mathbf{B}$, hence every element of V ; the solution space of $\chi_M(\partial_\tau)v = 0$ is exactly the stated quasi-exponential space. Conversely, differentiation maps each complete chain into itself because $\partial_\tau(\tau^k e^{\lambda\tau}) = (k\tau^{k-1} + \lambda\tau^k) e^{\lambda\tau}$; downward closure is essential (for example, $\text{span}\{\tau e^{-\tau}\}$ is not invariant). (b) Standard real Jordan theory. For (4): $\mathbf{1}$ and J commute and $J^2 = -\mathbf{1}$, so $e^{C\tau} = e^{\mu\tau} e^{\omega\tau J}$, and splitting the series of $e^{\omega\tau J}$ into even and odd powers gives $\cos(\omega\tau)\mathbf{1} + \sin(\omega\tau)J = R(\omega\tau)$. (c) Termwise integration of the exponential series; the two displayed cases are the semisimple-invertible and nilpotent extremes and combine by primary decomposition and direct sum. (d) Immediate. $\square$

Lemma 2.4 (Spanning = controllability). $\text{span}\{\mathbf{B}(\tau) : \tau \geq 0\} = \text{span}\{M^k \mathbf{B}(0) : k \geq 0\}$ (the Krylov space of the pair $(M, \mathbf{B}(0))$). Hence the curve family spans $\mathbb{R}^n$ iff $(M, \mathbf{B}(0))$ is controllable; otherwise the model reduces to the Krylov subspace, which is the minimal realisation.

Proof. $\tau \mapsto \mathbf{B}(\tau)$ is analytic; its span equals the span of all derivatives at 0, which is the Krylov space, an M -invariant subspace containing $\mathbf{B}(0)$ and contained in any such. $\square$

Remark 2.5 (Complex spectrum is not an extension). Theorem 2.3 makes the conceptual point of this whole program precise: the class forced by shift-invariance has exponents ranging over $\mathbb{C}$ with no further restriction. Real-eigenvalue bases (Nelson–Siegel [37], Svensson, exponential factor models) are the real slice; the precession blocks (4) are the rest of a class that was always inside the arbitrage-free theory. The eigenmodes, being $\text{Im } e^{(-1/\kappa + i\omega_n)\theta}$ up to normalisation, sit in the complex slice. Complex spectrum is only a compact representation of real sine–cosine factor pairs; it does not make rates or security values complex-valued.

2.4 Arbitrage-free affine realisations

The class has now been described by the shapes it admits and the drift it requires; what remains is to say which finite-dimensional models belong to it. The following standing assumption fixes the realisations against which the model will be measured.

Assumption 2.6. The curve is $r_t(\tau) = \varphi_t(\tau) + \mathbf{B}(\tau)^\top \mathbf{X}_t$ with $\mathbf{B}, \mathbf{X}_t \in \mathbb{R}^n$, where φ is deterministic, $t \mapsto \varphi_t$ is C^1 into H_w , and $t \mapsto \varphi_t$ is continuous into $D(\partial_\tau)$ equipped with its graph norm. Moreover $\mathbf{B}' = M\mathbf{B}$ with $(M, \mathbf{B}(0))$ controllable (Lemma 2.4). The factors follow $d\mathbf{X}_t = K\mathbf{X}_t dt + \Sigma_t d\mathbf{W}_t$ with a d -dimensional Brownian motion and a *deterministic*, locally continuous $\Sigma_t \in \mathbb{R}^{n \times d}$; write $C_t := \Sigma_t \Sigma_t^\top$. The realisation is required to be *consistent*: (3) holds for every initial factor state $\mathbf{X}_0 = x \in \mathbb{R}^n$, so the drift identity is an identity in (x, τ) , not merely a property of one realised path. (For the infinite-dimensional analytic background to this finite-dimensional setting, see [8, 15].)

Proposition 2.7 (Characterisation). *Under Assumption 2.6, the model is consistent — (3) holds for every initial state — if and only if*

$$K = M^\top \quad \text{and} \quad \dot{\varphi}_t(\tau) = \varphi'_t(\tau) + q_t(\tau), \quad q_t(\tau) := \mathbf{B}(\tau)^\top C_t \mathbf{I}(\tau), \quad (5)$$

with $\mathbf{I}$ as in Theorem 2.3(c); when M is invertible, $q_t(\tau) = \mathbf{B}(\tau)^\top C_t M^{-1}(\mathbf{B}(\tau) - \mathbf{B}(0))$. The transport equation for φ has the explicit solution

$$\varphi_t(\tau) = \varphi_0(\tau + t) + \int_0^t q_s(\tau + t - s) ds. \quad (6)$$

Proof. Volatility matching in (2) gives $\sigma_t(\tau) = \mathbf{B}(\tau)^\top \Sigma_t$, whence

$$\alpha_t(\tau) = \mathbf{B}(\tau)^\top \Sigma_t \int_0^\tau \Sigma_t^\top \mathbf{B}(u) du = \mathbf{B}(\tau)^\top C_t \mathbf{I}(\tau) = q_t(\tau).$$

The model's actual drift is $\dot{\varphi}_t(\tau) + \mathbf{B}(\tau)^\top K \mathbf{X}_t$, while (2) requires

$$\partial_\tau r_t(\tau) + \alpha_t(\tau) = \varphi'_t(\tau) + \mathbf{B}(\tau)^\top M^\top \mathbf{X}_t + q_t(\tau),$$

using $\partial_\tau(\mathbf{B}^\top \mathbf{X}) = (M\mathbf{B})^\top \mathbf{X} = \mathbf{B}^\top M^\top \mathbf{X}$. Now equate the two for every initial state and all τ . (At $t = 0$ the state is $\mathbf{X}_0 = x$; at $t > 0$ one has $\mathbf{X}_t = e^{Kt}x + \int_0^t e^{K(t-s)}\Sigma_s d\mathbf{W}_s$ with e^{Kt} invertible, so for each fixed noise path $\mathbf{X}_t$ still sweeps all of $\mathbb{R}^n$ as x does. Consistency therefore makes the identity hold for every $x \in \mathbb{R}^n$.) The $\mathbf{X}$ -linear parts give $\mathbf{B}(\tau)^\top(K - M^\top) = 0$ for all τ , hence $K = M^\top$ by Lemma 2.4; the deterministic parts give the transport equation. For (6), set $g_t(x) := \varphi_t(x - t)$; then $\dot{g}_t(x) = \dot{\varphi}_t - \varphi'_t = q_t(x - t)$, integrate, and substitute back. Note q_t must be deterministic for φ to absorb it, which is where deterministic Σ_t enters; a state-dependent Σ exits the class (cf. [35] for the framework; §6.1 constructs the affine relaxation). $\square$

Remark 2.8 (Support versus consistency; a counterexample). The quantifier in Proposition 2.7 matters. If (3) is required only along the realised process, the $\mathbf{X}$ -linear identity is tested only on the support of $\mathbf{X}_t$, and controllability of $(M, \mathbf{B}(0))$ does not make that support full. Take

$$\mathbf{B}(\tau) = \begin{pmatrix} e^{-\tau} \\ e^{-2\tau} \end{pmatrix}, \quad M = \text{diag}(-1, -2), \quad \Sigma = \begin{pmatrix} \sigma \\ 0 \end{pmatrix}, \quad K = \text{diag}(-1, 7), \quad \mathbf{X}_0 = 0.$$

Then $X_{2,t} \equiv 0$, so $\mathbf{B}^\top K \mathbf{X} = \mathbf{B}'^\top \mathbf{X}$ along the support although $K \neq M^\top$; taking $q(\tau) = \sigma^2(e^{-\tau} - e^{-2\tau})$ and φ the solution (6) of $\dot{\varphi}_t = \varphi'_t + q$, the pathwise drift condition holds. There are three ways to make the identification $K = M^\top$ rigorous. One may require consistency for every initial state, as in Proposition 2.7 and in the standard finite-dimensional-realisation formulation [6, 19]; one may assume the controllability Gramian of (K, Σ) is positive definite, which gives full support at every $t > 0$; or one may state the conclusion only on the reachable subspace. All completions constructed in this paper also satisfy the second condition: per block, $\{\Sigma\text{-column}, K\Sigma\text{-column}\}$ spans whenever $\omega_n \neq 0$, so for them the pathwise and consistent formulations coincide at positive times.

Remark 2.9 (Ho–Lee benchmark). $n = 1$, $b \equiv 1$, $M = 0$, $\Sigma_t = \sigma$: then $\mathbf{I}(\tau) = \tau$, $q(\tau) = \sigma^2\tau$, and (6) gives $\varphi_t(\tau) = \varphi_0(\tau + t) + \sigma^2 \int_0^t (\tau + t - s) ds = \varphi_0(\tau + t) + \sigma^2(t\tau + \frac{t^2}{2})$, which is precisely the Ho–Lee forward curve $f(t, T) = f(0, T) + \sigma^2 t(T - \frac{t}{2}) + \sigma W_t$ in Musiela coordinates. The zero eigenvalue illustrates Theorem 2.3(c): integration on $\ker M$ raises polynomial degree, and the convexity accumulates quadratically in the deterministic sector. This is an algebraic sanity check on a compact maturity interval (or in a larger state space), not an instance of Assumption 2.6 on the present infinite-horizon H_w . Although $b \equiv 1 \in H_w$, the convexity $q(\tau) = \sigma^2\tau$ and the resulting nonconstant polynomial deterministic sector are not in H_w .

Remark 2.10 ($\mathbb{P}$ versus $\mathbb{Q}$; essentially affine prices of risk). (5) constrains the $\mathbb{Q}$ -drift. Under $\mathbb{P}$,

$$d\mathbf{X}_t = (M^\top \mathbf{X}_t + \Sigma_t \gamma_t) dt + \Sigma_t d\mathbf{W}_t^\mathbb{P}.$$

If γ is constant (completely affine), the $\mathbb{P}$ - and $\mathbb{Q}$ -drift matrices coincide, although their deterministic drift intercepts may differ, and every spectral statement below applies to historical dynamics as well. If $\gamma_t = \gamma_0 + \Gamma_1 \mathbf{X}_t$ (essentially affine, empirically the relevant case [17]), then

$$K_t^\mathbb{P} = M^\top + \Sigma_t \Gamma_1.$$

When $\Sigma_t \equiv \Sigma$ is constant this is a fixed matrix $K^\mathbb{P} = M^\top + \Sigma \Gamma_1$, and the tests below restrict the distortion $\Sigma \Gamma_1$ (and Γ_1 only to the extent that it is identified through Σ). This matters because the source model is estimated under $\mathbb{P}$ [13].

Remark 2.11 (Strict realisations and exponent doubling). If one additionally demands that φ_t itself stay in a finite-dimensional family (a strict finite-dimensional realisation in the sense of [7]), then by Theorem 2.3(d) the family must contain the products entering q : exponents close under addition. For complex pairs $-\lambda_j \pm i\omega_j$ this generates *beat frequencies* $\omega_j \pm \omega_k$ at decay $\lambda_j + \lambda_k$: convexity lives at sums and differences of the basis frequencies. Because integration also introduces a zero-exponent term, q can retain the original exponents as well as their pairwise sums. On the model's lattice $\omega_n = n\pi/\theta^*$ the beat frequencies remain on the lattice, although their decay exponents need not coincide with those of the original loadings. This is a small structural felicity of the equally spaced spectrum.

The yardstick is assembled: the shapes the class admits, the drift it requires, and the realisations that satisfy both. What remains is to place the model against it.

3 The cylindrical SPDE, rigorously

The model enters here, and so does the reason the comparison cannot be made pointwise: under cylindrical noise the field is too rough for the drift condition to be posed maturity-by-maturity. This section makes both statements exact — the model well-posed, the pointwise route closed — which is what forces the diagnosis onto finite-dimensional invariant subspaces, where the next section conducts it.

3.1 Model and spectral data

Throughout, $\theta^* > 0$, $\kappa > 0$ and $\sigma_0 > 0$; these standing assumptions are used from here on without further mention. They are load-bearing rather than cosmetic. With $\kappa < 0$ the sign of every λ_n reverses and (7) becomes a backward heat equation; with $\sigma_0 = 0$ there is no noise, and both the variance bound of Proposition 3.3(a) and the nondegeneracy step of Theorem 4.6 use $\sigma_0 > 0$.

On $[0, \theta^*]$ (after translating $\theta_{\min}$ to 0), the deformation field solves

$$dX_t = \underbrace{\left[\partial_\theta + \frac{\kappa}{2} \partial_\theta^2 \right]}_{=: \tilde{A}} X_t dt + \sigma_0 d\mathcal{B}_t, \quad X_t(0) = X_t(\theta^*) = 0, \quad (7)$$

where $\mathcal{B}$ is a cylindrical Brownian motion on $H = L^2([0, \theta^*], e^{2\theta/\kappa} d\theta)$: a family of centred Gaussian random functionals with $\mathbb{E}[\mathcal{B}_t(\phi)\mathcal{B}_s(\psi)] = (t \wedge s) \langle \phi, \psi \rangle_H$ [15]. Two objects must be kept apart throughout: the *formal* differential expression $\tilde{A} := \partial_\theta + \frac{\kappa}{2} \partial_\theta^2$, which acts on any smooth function, and its *closed realization* A in H on the Dirichlet domain $D(A) = \{h \in H^2 : h(0) = h(\theta^*) = 0\}$, which generates the semigroup and the mild solution. The distinction is not pedantry — it is exactly where the model's difficulty lives, since the shift-invariant envelope of §4 leaves $D(A)$ while remaining inside the domain of $\tilde{A}$. The closed operator A has orthonormal eigenbasis

$$Ae_n = -\lambda_n e_n, \quad e_n(\theta) = \sqrt{\frac{2}{\theta^*}} \sin(\omega_n \theta) e^{-\theta/\kappa}, \quad \omega_n = \frac{n\pi}{\theta^*}, \quad \lambda_n = \frac{1}{2\kappa} (1 + \kappa^2 \omega_n^2), \quad (8)$$

and coordinates $x_n(t) = \langle X_t, e_n \rangle_H$ are independent OU processes $dx_n = -\lambda_n x_n dt + \sigma_0 dW_t^n$.

Proposition 3.1 (Symbol). *For $z \in \mathbb{C}$, $\tilde{A}e^{z\theta} = p(z)e^{z\theta}$ with $p(z) = z + \frac{\kappa}{2}z^2$. This is an identity for the formal expression: apart from the sine combinations selected below, these characters are not in $D(A)$. The map $z \mapsto p(z)$ is real-valued exactly on $\mathbb{R} \cup \{\operatorname{Re} z = -1/\kappa\}$; on the critical line $-\frac{1}{\kappa} + i\omega$ it takes the value $-\lambda(\omega) = -\frac{1+\kappa^2\omega^2}{2\kappa}$. The closed Dirichlet operator has the real spectrum $\operatorname{spec}(A) = \{-\lambda_n : n \geq 1\}$. For each operator eigenvalue $-\lambda_n$, the spatial characteristic equation $p(z) = -\lambda_n$ has the conjugate roots $-1/\kappa \pm i\omega_n$. These characteristic exponents lie on the critical line, the axis of symmetry of the parabola p ; equivalently, they form the spectrum of the maturity-generator block M_n introduced in Proposition 4.1.*

Proof. Termwise, $\tilde{A}e^{z\theta} = (z + \frac{\kappa}{2}z^2)e^{z\theta} = p(z)e^{z\theta}$. For the reality locus write $z = x + iy$: $\operatorname{Im} p(z) = y + \frac{\kappa}{2} \cdot 2xy = y(1 + \kappa x)$, which vanishes iff $y = 0$ or $x = -1/\kappa$. Centering on the critical line, $z = -\frac{1}{\kappa} + w$ gives

$$p(-\frac{1}{\kappa} + w) = -\frac{1}{\kappa} + w + \frac{\kappa}{2} \left(\frac{1}{\kappa^2} - \frac{2w}{\kappa} + w^2 \right) = -\frac{1}{2\kappa} + \frac{\kappa}{2}w^2,$$

even in w : the line is the parabola's axis of symmetry. At $w = i\omega$ the value is $-\frac{1+\kappa^2\omega^2}{2\kappa} = -\lambda(\omega)$, and $p(z) = -\lambda_n$ becomes $\frac{\kappa}{2}w^2 = -\frac{\kappa}{2}\omega_n^2$, i.e. $w = \pm i\omega_n$: the conjugate pair $-\frac{1}{\kappa} \pm i\omega_n$. The Dirichlet conditions select from each root pair the combination $\operatorname{Im} e^{(-1/\kappa + i\omega)\theta} \propto e^{-\theta/\kappa} \sin \omega\theta$ vanishing at 0, and vanishing at θ^* quantises ω to $\omega_n = n\pi/\theta^*$ — which is (8). $\square$

This is the reason the eigenfunctions are damped sines with a *common* damping rate. The imaginary part of p vanishes on the critical line, so both members of each conjugate pair — sine and cosine alike — are formal eigenfunctions of $\tilde{A}$ with the same real eigenvalue; only the sine combination satisfies the Dirichlet conditions and so lies in $D(A)$. That observation does the heavy lifting in §4.

3.2 Well-posedness under cylindrical noise

Before the model can be compared with anything it has to exist as a function-valued process, and under cylindrical noise that is not automatic. The operator that secures it is the one that will later break the drift condition.

Proposition 3.2 (Function-valued solutions). *For deterministic $X_0 \in H$, the mild solution*

$$X_t = e^{tA}X_0 + \sigma_0 \int_0^t e^{(t-s)A} d\mathcal{B}_s$$

of (7) is a well-defined H -valued Gaussian process with a continuous modification, and the stationary variance $\operatorname{Var} X_\infty(\theta) = \sigma_0^2 \sum_n \frac{e_n(\theta)^2}{2\lambda_n}$ is finite pointwise and in H . The second-order term is essential: with the shift generator alone, the semigroup is not Hilbert–Schmidt-improving and the stochastic convolution with cylindrical noise has no H -valued version [15]. (This is a different operator, not the singular substitution $\kappa = 0$ in the eigenfunctions (8).) For an arbitrary H -valued initial condition independent of the future noise, the same existence and path-continuity conclusions hold, but the solution need not be Gaussian.

Proof. Diagonalise. From (8), $e^{sA}e_n = e^{-\lambda_n s}e_n$, so the semigroup is Hilbert–Schmidt on H for every $s > 0$, with $\|e^{sA}\|_{\text{HS}}^2 = \sum_n e^{-2\lambda_n s}$ and

$$\int_0^t \|e^{sA}\|_{\text{HS}}^2 ds = \sum_n \frac{1 - e^{-2\lambda_n t}}{2\lambda_n} \leq \frac{1}{2} \sum_n \lambda_n^{-1} < \infty.$$

By the theory of linear equations with additive noise [15, Ch. 5], the stochastic convolution is therefore a well-defined centred H -valued Gaussian process with trace-class covariance $Q_t = \sigma_0^2 \sum_n \frac{1-e^{-2\lambda_n t}}{2\lambda_n} e_n \otimes e_n$; adding $e^{tA} X_0$ gives the mild solution. It has a continuous H -valued modification by the factorization method [15, §5.3]: for $0 < \alpha < \frac{1}{4}$,

$$\int_0^T s^{-2\alpha} \|e^{sA}\|_{\text{HS}}^2 ds \leq \Gamma(1-2\alpha) \sum_n (2\lambda_n)^{2\alpha-1} < \infty,$$

since $\lambda_n \asymp n^2$ and $2(2\alpha-1) < -1$ — the same exponent $\frac{1}{4}$ that resurfaces as the temporal Hölder threshold in Proposition 3.3(b). In coordinates the solution is the OU family below (8); the stationary variances $\sigma_0^2/(2\lambda_n)$ are summable, and $\sup_\theta e_n(\theta)^2 \leq 2/\theta^*$ makes $\sigma_0^2 \sum_n e_n(\theta)^2/(2\lambda_n)$ converge uniformly on $[0, \theta^*]$: the stationary field exists pointwise and in H .

For the pure shift, drop the diffusion term. The generator ∂_θ (a first-order operator admits one endpoint condition; take the one at θ^*) generates the killed shift $(T_s \phi)(\theta) = \phi(\theta + s) \mathbf{1}_{\{\theta+s \leq \theta^*\}}$. In $H = L^2(e^{2\theta/\kappa} d\theta)$ its adjoint is $(T_s^* \psi)(\theta) = e^{-2s/\kappa} \psi(\theta - s) \mathbf{1}_{\{\theta \geq s\}}$, so $T_s T_s^*$ is multiplication by $e^{-2s/\kappa} \mathbf{1}_{[0, \theta^*-s]}$, a strictly positive multiple of a projection with infinite-dimensional range. Hence $\|T_s\|_{\text{HS}}^2 = \text{tr}(T_s T_s^*) = \infty$ for every $s < \theta^*$: no smoothing, no Hilbert–Schmidt improvement. The would-be covariance of the stochastic convolution is multiplication by $g_t(\theta) = \sigma_0^2 \frac{\kappa}{2} (1 - e^{-2(t \wedge (\theta^* - \theta))/\kappa})$, which is bounded below on sets of positive measure. A multiplication operator with such a multiplier has infinite trace, since its quadratic form dominates a positive multiple of an infinite-rank projection. No H -valued Gaussian vector has this covariance, so the convolution admits no H -valued version [15, Ch. 5]: transported without smoothing, cylindrical noise stays cylindrical. $\square$

So the local-deformation operator $\frac{\kappa}{2} \partial_\theta^2$ is what makes “one source of randomness per maturity” mathematically possible at all. The same operator will turn out to be exactly what breaks the drift condition (Theorem 4.3); that is the precise form of the tension in [13].

3.3 Pointwise regularity: forward rates are not semimartingales

Existence in H does not make individual forward rates tractable. The next result closes the pointwise route: at a fixed maturity the rate is too rough for the drift condition of Lemma 2.1 to be posed at all, which is what sends the comparison onto finite-dimensional invariant subspaces.

Proposition 3.3. *Let X solve (7) either from its stationary law or from a deterministic $X_0 \in H$, and fix $\theta \in (0, \theta^*)$. In the stationary case let $I = [0, T]$, while in the deterministic-initial case let $I = [t_*, t_* + T] \subset (0, \infty)$. Then:*

- (a) *uniformly for $t, t+h \in I$ and sufficiently small $h > 0$, there are $0 < c_{\theta, I} \leq C_{\theta, I} < \infty$ such that*

$$c_{\theta, I} \sqrt{h} \leq \text{Var}(X_{t+h}(\theta) - X_t(\theta)) \leq C_{\theta, I} \sqrt{h};$$

- (b) *$t \mapsto X_t(\theta)$ has a modification that is locally Hölder of every order $< \frac{1}{4}$ on I ;*

- (c) *along the uniform partitions of I into m subintervals, the quadratic-variation sums Q_m satisfy $\mathbb{E} Q_m \asymp \sqrt{m} \rightarrow \infty$ and, for all sufficiently large m ,*

$$\mathbb{P}(Q_m \geq c_I \sqrt{m}) \geq \frac{1}{12}$$

for some $c_I > 0$. Hence (Q_m) is not tight and cannot converge in probability to a finite limit; in particular, $t \mapsto X_t(\theta)$ is not a semimartingale.

Cf. [41, 40] for the analogous statements for the stochastic heat equation, including the finiteness of the quartic variation, of which (a)–(c) are the drift-augmented transcription.

Proof. (a) $\text{Var}(X_{t+h}(\theta) - X_t(\theta)) = \sigma_0^2 \sum_n \lambda_n^{-1} (1 - e^{-\lambda_n h}) e_n(\theta)^2$ by the OU covariances. With $e_n(\theta)^2 = \frac{2}{\theta^*} \sin^2(\omega_n \theta) e^{-2\theta/\kappa} \leq \frac{2}{\theta^*}$ and $\lambda_n \asymp n^2$, split at $n_* = \lceil h^{-1/2} \rceil$: the sum is $\lesssim \sum_{n \leq n_*} h + \sum_{n > n_*} n^{-2} \asymp \sqrt{h}$. For the lower bound, $\sin^2(n\pi\theta/\theta^*)$ has positive lower density bounded away from zero (for irrational θ/θ^* by Weyl equidistribution with mean $\frac{1}{2}$; for rational p/q the values are q -periodic and vanish only when $q \mid n$), so a positive-density subsequence contributes $\gtrsim \sqrt{h}$. (b) Gaussianity plus (a) gives $\mathbb{E}|X_{t+h} - X_t|^{2p} \leq C_p h^{p/2}$ and Kolmogorov's criterion yields Hölder $< \frac{1}{4}$. (c) $\mathbb{E}Q_m = \sum_i \text{Var}(\Delta_i X) \asymp m \cdot \sqrt{T/m} = \sqrt{Tm}$. For jointly Gaussian increments, $\text{Var} Q_m = 2 \sum_{i,j} \text{Cov}(\Delta_i, \Delta_j)^2 \leq 2 (\sum_i \text{Var} \Delta_i)^2 = 2(\mathbb{E}Q_m)^2$, so $\mathbb{E}Q_m^2 \leq 3(\mathbb{E}Q_m)^2$ and Paley–Zygmund gives $\mathbb{P}(Q_m \geq \frac{1}{2} \mathbb{E}Q_m) \geq \frac{(\mathbb{E}Q_m)^2}{4\mathbb{E}Q_m^2} \geq \frac{1}{12}$. A continuous semimartingale has Q_m converging in probability to its finite quadratic variation and hence has a tight sequence of quadratic-variation sums. Here, for every $R < \infty$ and all sufficiently large m ,

$$\mathbb{P}(Q_m > R) \geq \mathbb{P}(Q_m \geq \frac{1}{2} \mathbb{E}Q_m) \geq \frac{1}{12},$$

because $\mathbb{E}Q_m \rightarrow \infty$. Thus (Q_m) is not tight, a contradiction.

Deterministic initial data. Write $X_t = e^{tA} X_0 + Z_t$ with Z the stochastic convolution, which for $t > 0$ is defined pointwise (e^{tA} is analytic, so $e^{tA} X_0$ is smooth and deterministic). Only Z is random, and its modal variances are $v_n(t) = \frac{\sigma_0^2}{2\lambda_n} (1 - e^{-2\lambda_n t})$, giving the exact increment law

$$\text{Var}(Z_{t+h}(\theta) - Z_t(\theta)) = \sum_n e_n(\theta)^2 \left[v_n(t) (1 - e^{-\lambda_n h})^2 + \frac{\sigma_0^2}{2\lambda_n} (1 - e^{-2\lambda_n h}) \right].$$

The second bracket alone reproduces the two-sided $\sqrt{h}$ estimate of (a), uniformly for $t \in I$, by the same split at $n_* = \lceil h^{-1/2} \rceil$. The first is nonnegative and bounded by $\frac{\sigma_0^2}{2\lambda_n} (1 - e^{-\lambda_n h})^2$, which the same split bounds by $O(\sqrt{h})$. Thus the preceding Gaussian argument applies to Z on every $I = [t_*, t_* + T] \subset (0, \infty)$: if Q_m^Z denotes its quadratic-variation sum, then $\mathbb{E}Q_m^Z \asymp \sqrt{m}$ and

$$\mathbb{P}\left(Q_m^Z \geq \frac{1}{2} \mathbb{E}Q_m^Z\right) \geq \frac{1}{12}.$$

Put $D_t := (e^{tA} X_0)(\theta)$. Analytic smoothing makes D continuously differentiable with bounded derivative on I , and therefore its quadratic-variation sum satisfies

$$Q_m^D := \sum_i (\Delta_i D)^2 \leq \frac{T^2}{m} \sup_{s \in I} |D'_s|^2 = O(m^{-1}).$$

Regarding the increment vectors as elements of $\mathbb{R}^m$, the triangle inequality gives

$$\sqrt{Q_m^X} \geq \sqrt{Q_m^Z} - \sqrt{Q_m^D}.$$

Consequently, on the preceding Paley–Zygmund event and for all sufficiently large m , $Q_m^X \geq \frac{1}{4}Q_m^Z \gtrsim \sqrt{m}$. Moreover, since the increments of Z are centred and D is deterministic,

$$\mathbb{E}Q_m^X = \mathbb{E}Q_m^Z + Q_m^D \asymp \sqrt{m}.$$

This proves (b)–(c) for deterministic initial data without requiring a pointwise value of X_0 at time zero. $\square$

Remark 3.4 (Where the drift condition can live). Proposition 3.3 means the pointwise HJM computation of Lemma 2.1 is not even well-posed for (7): individual forward rates are not Itô processes. Pairings $\langle X_t, \phi \rangle$ with $\phi \in D(A^*)$ are semimartingales (weak formulation). On the finite maturity interval the bond integral $X \mapsto \int_0^u X(\theta) d\theta$ is a bounded functional on $H = L^2([0, \theta^*], e^{2\theta/\kappa} d\theta)$, with H -representer $e^{-2\theta/\kappa} \mathbf{1}_{[0, u]}(\theta)$. This representer is *not* in $D(A^*)$: under the unitary conjugacy of Proposition 4.10 it becomes $e^{-\theta/\kappa} \mathbf{1}_{[0, u]}$, which jumps at u and so is not even H^1 . That alone does not decide the question that matters — whether traded-bond processes are semimartingales under the fully infinite cylindrical dynamics — which remains open (Proposition 4.10 and the discussion closing §4). The requirement is not special to this model: a state space in which bond prices are to be Itô processes must either make the maturity functionals that define them semimartingale pairings or carry the traded coordinates explicitly. That is the division of labour between the two spaces used here: the arbitrage-free class of §2 lives in H_w , where point evaluation is continuous, while the SPDE itself is posed in H . At every finite truncation the bond processes are ordinary Itô semimartingales, and Theorem 4.6 below constructs an explicit finite-bond witness. The random-field literature [33, 22, 39] handles this with spatially *correlated* noise, for which a drift condition in terms of the covariance kernel exists; white-in-maturity noise is exactly the singular case. The rigorous route is therefore the one taken in §§4–5: diagnose and repair on finite-dimensional invariant subspaces, then control the limit — where the singularity resurfaces in a precisely quantified form (Theorem 5.2).

4 The obstruction: symbol, witness, and magnitude

The comparison between the class of §2 and the model of §3 now runs at three increasing levels of concreteness. In factor coordinates, a symbol computation identifies exactly which dynamics the class tolerates and shows that none of the model’s blocks is among them. In traded securities, the failure becomes an explicit arbitrage portfolio (§4.4). In magnitude, a normalized quotient shows that refinement does not attenuate it (§4.5). §4.6 asks what survives in the infinite model, and §4.7 what changes when the silent coordinates are given noise of their own: the impossibility becomes a price.

4.1 The invariant envelope doubles

The comparison needs a shift-invariant family, and the model’s sine modes do not form one. The first step is therefore to find the smallest invariant space that contains them.

Proposition 4.1 (Envelope doubling). *Let*

$$V_n := \text{span}\{e^{-\theta/\kappa} \sin \omega_n \theta, e^{-\theta/\kappa} \cos \omega_n \theta\}.$$

Then V_n is the minimal ∂_θ -invariant subspace containing e_n , with generating matrix

$$M_n = \begin{pmatrix} -1/\kappa & \omega_n \\ -\omega_n & -1/\kappa \end{pmatrix}, \quad \text{spec } M_n = \left\{ -\frac{1}{\kappa} \pm i\omega_n \right\}, \quad \det M_n = \frac{1}{\kappa^2} + \omega_n^2. \quad (9)$$

Consequently the minimal shift-invariant envelope of the N -mode truncation is

$$\text{span}\{e_1, \dots, e_N\} \mapsto \bigoplus_{n \leq N} V_n,$$

of dimension $2N$. The cosine partners violate the Dirichlet conditions (cos-modes are nonzero at $\theta = 0$), which is why the model's state space itself is not shift-invariant.

Economically, the cosine partner is not an additional source of randomness. It is the deterministic roll-down direction generated as a sine loading ages; omitting it prevents the finite state from representing maturity transport.

Proof. Differentiating the pair of characters,

$$\partial_\theta e^{-\theta/\kappa} \begin{pmatrix} \sin \omega_n \theta \\ \cos \omega_n \theta \end{pmatrix} = e^{-\theta/\kappa} \begin{pmatrix} -1/\kappa & \omega_n \\ -\omega_n & -1/\kappa \end{pmatrix} \begin{pmatrix} \sin \omega_n \theta \\ \cos \omega_n \theta \end{pmatrix},$$

which is (9). V_n is 2-dimensional and invariant; no 1-dimensional invariant subspace contains e_n because a 1-dimensional invariant subspace is spanned by a real eigenvector of M_n , and $\text{spec } M_n \cap \mathbb{R} = \emptyset$. Distinct n give linearly independent trigonometric characters after dividing by the common exponential, so the envelopes sum directly. $\square$

The envelope is the object on which the comparison is made from here on. That it leaves the Dirichlet domain is the first thing the next subsection has to settle.

4.2 The generator comparison is exact on the envelope

The domain distinction of §3 decides which object the comparison is about. The cosine partners adjoined in Proposition 4.1 violate the Dirichlet conditions and do not belong to $D(A)$; on the envelope the relevant operator is therefore the formal expression $\tilde{A} = \partial_\theta + \frac{\kappa}{2} \partial_\theta^2$, acting on the analytic (quasi-exponential) extensions. The obstruction proved here is therefore a *tangency/symbol* statement, in the spirit of the invariant-manifold calculus of [21]: the formal symbol acts one way on the invariant envelope and no-arbitrage transport demands another. It concerns the formal expression, not an invariant subspace of the closed Dirichlet operator. The sine truncation itself is invariant under A ; its minimal shift-invariant envelope leaves $D(A)$ (Proposition 4.1), and the completions built on that envelope are finite-dimensional models in their own right.

Proposition 4.2 (Induced factor drifts: the symbol of a curve dynamics). *Let V be any finite-dimensional shift-invariant space with basis $\mathbf{B}$, $\mathbf{B}' = M\mathbf{B}$, and consider curve dynamics $d(\mathbf{B}^\top \mathbf{X}_t) = (G\mathbf{B}^\top \mathbf{X}_t) dt + (\text{noise})$ driven by a generator G leaving V invariant.*

- (a) (Any constant-coefficient generator.) *Let $G = \sum_{k=0}^d a_k \partial_\theta^k$ with symbol $p(z) = \sum_{k=0}^d a_k z^k$. Then V is G -invariant and the induced factor drift matrix is*

$$K^{\text{def}} = p(M)^\top = p(M^\top).$$

The deformation generator of this paper is the case $G = \tilde{A} = \partial_\theta + \frac{\kappa}{2} \partial_\theta^2$ (the model's formal expression, per the caveat above), with $p(z) = z + \frac{\kappa}{2} z^2$.

(b) If $G = \partial_\tau$ with the affine inhomogeneity of (2) (arbitrage-free), the required factor drift is $K^{\text{NA}} = M^\top$ (Proposition 2.7).

(c) On the model's blocks (9), $p(M_n) = -\lambda_n \mathbf{1}_2$ exactly, consistent with Proposition 3.1: the symbol is real on the critical line, so A acts isotropically there — decay λ_n , no rotation.

Proof. (a) $\mathbf{B}' = M\mathbf{B}$ iterates to $\mathbf{B}^{(k)} = M^k \mathbf{B}$, so $G\mathbf{B} = \sum_k a_k \mathbf{B}^{(k)} = p(M)\mathbf{B}$; in particular V is G -invariant. For a curve $\mathbf{B}^\top x$, $G(\mathbf{B}^\top x) = (p(M)\mathbf{B})^\top x = \mathbf{B}^\top (p(M)^\top x)$, and since the components of $\mathbf{B}$ are a basis, coordinate representations are unique: the curve drift $G\mathbf{B}^\top \mathbf{X}_t$ is carried by the factor drift matrix $K^{\text{def}} = p(M)^\top$, which equals $p(M^\top)$ because transposition fixes polynomials in a single matrix. (b) is Proposition 2.7: under the arbitrage-free dynamics the affine inhomogeneity is absorbed by φ , and the factor block of ∂_τ is $M^\top$. (c) With $M_n = -\frac{1}{\kappa} \mathbf{1} + \omega_n J$ and $J^2 = -\mathbf{1}$, $M_n^2 = \left(\frac{1}{\kappa^2} - \omega_n^2\right) \mathbf{1} - \frac{2\omega_n}{\kappa} J$, so

$$p(M_n) = M_n + \frac{\kappa}{2} M_n^2 = \left(-\frac{1}{\kappa} + \frac{\kappa}{2} \left(\frac{1}{\kappa^2} - \omega_n^2\right)\right) \mathbf{1} + \left(\omega_n - \frac{\kappa}{2} \cdot \frac{2\omega_n}{\kappa}\right) J = -\lambda_n \mathbf{1} :$$

the J -part cancels identically — the algebraic form of “real symbol on the critical line” (Proposition 3.1): decay λ_n , no rotation. $\square$

Theorem 4.3 (The symbol test; the exact obstruction). *With the notation of Proposition 4.2, let V be a shift-invariant block with generating matrix M and let the curve dynamics have symbol p .*

(a) (The test, for any symbol.) *The dynamics carries the arbitrage-free factor drift on V — that is, $K^{\text{def}} = K^{\text{NA}}$ — if and only if*

$$\boxed{m_M(z) \mid p(z) - z}$$

where m_M is the minimal polynomial of M . When M is diagonalisable this says exactly that the symbol must fix the block's spectrum pointwise, $p(\lambda) = \lambda$ for every $\lambda \in \text{spec } M$; in general $p(z) - z$ must vanish at each eigenvalue to the order of its largest Jordan block. Thus consistency of the factor drift on each finite-dimensional shift-invariant block is reduced to one polynomial divisibility; on the envelope $\bigoplus_n V_n$ the induced drift is block-diagonal, so the condition holds there if and only if it holds on every block.

(b) (The instance.) For $p(z) = z + \frac{\kappa}{2} z^2$ the criterion reads $m_M \mid \frac{\kappa}{2} z^2$, i.e.

$$K^{\text{def}} = K^{\text{NA}} \iff p(M) = M \iff M^2 = 0 \iff V \subseteq \ker \partial_\theta^2 = \{\text{affine curves } a + b\theta\}.$$

Since $\text{spec } M \subset \{\text{Re } z = -1/\kappa\}$ contains no zero eigenvalue, no block of this model passes: the drift condition fails on every genuinely decaying block ($\mu \neq 0$) and every oscillatory block ($\omega \neq 0$). The quantitative discrepancy is $p(M) - M = \frac{\kappa}{2} M^2$, i.e. on the spectrum, $p(z) - z = \frac{\kappa}{2} z^2$. In H_w the affine sliver is thinner still: since $w \geq 1$ is nondecreasing, $\int_0^\infty w = \infty$ and $\|a + b\theta\|_w < \infty$ forces $b = 0$, so inside the state space the tolerated class is the constants alone. For the model's blocks the two induced dynamics are

$$K_n^{\text{def}} = -\lambda_n \mathbf{1}_2 \quad \text{vs} \quad K_n^{\text{NA}} = M_n^\top = \begin{pmatrix} -1/\kappa & -\omega_n \\ \omega_n & -1/\kappa \end{pmatrix},$$

differing in exactly two ways: the decay mismatch $\kappa\lambda_n = \frac{1}{2}(1 + \kappa^2\omega_n^2)$, equal to 1 only if the lattice contains an index satisfying $\kappa\omega_n = 1$. Such an index exists precisely when $\theta^*/(\pi\kappa) \in \mathbb{N}$; otherwise $\theta^*/(\pi\kappa)$ is only the continuous crossing location. The second discrepancy is the missing rotation $\pm\omega_n$. Separately and non-spectrally, the convexity q_t of (5) is absent from (7).

Proof. (a) $K^{\text{def}} = p(M)^\top$ and $K^{\text{NA}} = M^\top$; transposing, equality is $p(M) = M$, i.e. $q(M) = 0$ for $q := p - \text{id}$. A polynomial annihilates M precisely when it is divisible by the minimal polynomial, which is the stated criterion; the spectral form is the standard reading of $m_M \mid q$ (simple roots when M is diagonalisable, root multiplicities equal to the largest Jordan-block sizes in general). (b) Here $q = \frac{\kappa}{2}z^2$, so $q(M) = 0$ reads $\frac{\kappa}{2}M^2 = 0$. On a 2×2 rotation-decay cell $M = \mu\mathbf{1} + \omega J$, $M^2 = (\mu^2 - \omega^2)\mathbf{1} + 2\mu\omega J = 0$ forces $\mu = \omega = 0$; a nilpotent Jordan cell $\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ also satisfies $M^2 = 0$ and spans $\{1, \theta\}$. Since $\mathbf{B}'' = M^2\mathbf{B}$ and the components of $\mathbf{B}$ form a basis, $M^2 = 0$ is equivalent to $V \subseteq \ker \partial_\theta^2$. The displayed block pair is $K_n^{\text{def}} = p(M_n)^\top = -\lambda_n\mathbf{1}_2$ (Proposition 4.2(c)) against $K_n^{\text{NA}} = M_n^\top$ (Proposition 2.7); the mismatch ratio $\kappa\lambda_n = \frac{1}{2}(1 + \kappa^2\omega_n^2)$ restates (8). $\square$

Remark 4.4 (The obstruction is measure-invariant). The discrepancy of Theorem 4.3 is not a market price of risk. On the Brownian filtration every equivalent measure change is a Girsanov tilt $d\mathbf{W}^\mathbb{Q} = d\mathbf{W} + \gamma_t dt$, so it changes the factor drift by $-\Sigma\gamma_t$ and therefore can move it only within the range of the volatility. With the sine-only loading, range Σ is the sine coordinates, while per block

$$K_n^{\text{def}} - K_n^{\text{NA}} = -\lambda_n\mathbf{1}_2 - M_n^\top = \begin{pmatrix} \frac{1}{\kappa} - \lambda_n & \omega_n \\ -\omega_n & \frac{1}{\kappa} - \lambda_n \end{pmatrix}$$

has cosine row $(-\omega_n, \frac{1}{\kappa} - \lambda_n) \neq 0$: the missing rotation acts on the *noiseless* coordinate, which no tilt γ — adapted, state-dependent, or otherwise — can reach. (Even when a matched lattice index n^* with $\kappa\omega_{n^*} = 1$ exists, so the decay entry $\frac{1}{\kappa} - \lambda_n = -\frac{(\kappa\omega_n - 1)(\kappa\omega_n + 1)}{2\kappa}$ vanishes, the rotation entry $-\omega_{n^*}$ does not.) The missing convexity is likewise unreachable: q contains an original-frequency term at decay $1/\kappa$, together with a zero-frequency term and $2\omega_n$ beat terms at decay $2/\kappa$ (Remark 2.11); the latter components lie outside the sine-only volatility range. Consequently the envelope realisation equipped with the induced dynamics admits no equivalent local martingale measure, for *every* N : at each truncation order the verdict is not “wrong $\mathbb{Q}$ -drift” but “no $\mathbb{Q}$ exists.” This sharpens Remark 2.10: the $\mathbb{P}$ -freedom bought by an essentially affine price of risk is exactly $\Sigma\Gamma_1$, the sine rows, and the deviation exits it identically. And it is visible only on the doubled envelope; the cosine partners adjoined in Proposition 4.1 are what expose it. Theorem 4.6 below turns this factor-range statement into an explicit arbitrage using finitely many traded zero-coupon bonds.

4.3 Spectral locking and factor autocovariance

Theorem 4.3 was stated for the model. Its characterisation half, Proposition 2.7, has a consequence for every consistent model that is worth recording on its own, because it is testable and the model violates it maximally.

Corollary 4.5 (Spectral locking). *In a consistent arbitrage-free model of Assumption 2.6 (Proposition 2.7), $\text{spec } K = \text{spec } M$ with multiplicities. Moreover, if $\Sigma_t \equiv \Sigma$ is constant with $C = \Sigma\Sigma^\top$, $\text{Respec } M < 0$, and the factor process is initialised at (or has relaxed to) its stationary*

law, then the stationary covariance V_∞ solves the Lyapunov equation $M^\top V_\infty + V_\infty M + C = 0$ and the stationary autocovariance is

$$R(h) := \text{Cov}(\mathbf{X}_{t+h}, \mathbf{X}_t) = e^{M^\top h} V_\infty.$$

On a semisimple complex block with spectrum $-\lambda \pm i\omega$ this equals $e^{-\lambda h} R(\omega h)^\top V_\infty$ by (4): on every stochastically excited such block, the nonzero covariance entries of its rows are damped sinusoids with the maturity frequency ω and decay rate λ . On a complex Jordan chain of length $m > 1$ the frequency and the decay rate are the same, but the entries are $e^{-\lambda h}$ times sine-cosine combinations with polynomial coefficients of degree at most $m - 1$: it is the pure-sinusoid shape, not the spectral content, that requires semisimplicity. Thus a maturity ripple of wavelength $2\pi/\omega$ corresponds, on that block, to a calendar-time cycle of period $2\pi/\omega$.

Proof. Proposition 2.7 gives $K = M^\top$, which is similar to M , so the Jordan structure — in particular the spectrum with multiplicities — coincides. With constant Σ and $\text{Respec } M < 0$, the factor OU system $d\mathbf{X} = M^\top \mathbf{X} dt + \Sigma d\mathbf{W}$ has stationary law $N(0, V_\infty)$ with $V_\infty = \int_0^\infty e^{M^\top s} C e^{Ms} ds$; the integral converges, and

$$M^\top V_\infty + V_\infty M = \int_0^\infty \frac{d}{ds} (e^{M^\top s} C e^{Ms}) ds = -C,$$

which is the Lyapunov equation, uniquely solvable because $\text{spec } M^\top \cap \text{spec}(-M) = \emptyset$. For $h \geq 0$, $\mathbf{X}_{t+h} = e^{M^\top h} \mathbf{X}_t + \int_t^{t+h} e^{M^\top(t+h-u)} \Sigma d\mathbf{W}_u$ with the integral independent of $\mathcal{F}_t$, so under stationarity $R(h) = \text{Cov}(\mathbf{X}_{t+h}, \mathbf{X}_t) = e^{M^\top h} V_\infty$. Work in coordinates adapted to the real Jordan form (the completions' coordinates already are). On a semisimple cell $M_b = -\lambda \mathbf{1} + \omega J$, (4) gives the block rows $e^{M_b^\top h} = e^{-\lambda h} R(\omega h)^\top$, so for every column j

$$\begin{aligned} R(h)_{1j} &= e^{-\lambda h} (\cos(\omega h)(V_\infty)_{1j} - \sin(\omega h)(V_\infty)_{2j}), \\ R(h)_{2j} &= e^{-\lambda h} (\sin(\omega h)(V_\infty)_{1j} + \cos(\omega h)(V_\infty)_{2j}), \end{aligned}$$

i.e. $R(h)_{ij} = \rho_{ij} e^{-\lambda h} \cos(\omega h - \phi_{ij})$ with amplitude $\rho_{ij} = ((V_\infty)_{1j}^2 + (V_\infty)_{2j}^2)^{1/2}$: every nonvanishing entry in those rows is a damped sinusoid at frequency exactly ω and decay exactly λ . The block is stochastically excited — some $\rho_{ij} \neq 0$ — whenever its rows Σ_b of Σ are nonzero: the block-diagonal of V_∞ is then the controllability Gramian $\int_0^\infty e^{M_b^\top s} \Sigma_b \Sigma_b^\top e^{M_b s} ds$, positive definite because $\{v, (-\lambda \mathbf{1} - \omega J)v\}$ spans $\mathbb{R}^2$ for every $v \neq 0$ when $\omega \neq 0$ ($Jv \perp v$). For a chain of length m the real Jordan block is $M_b = (\mathbf{1}_m \otimes C) + (N \otimes \mathbf{1}_2)$ with N the nilpotent shift; the two summands commute, so $e^{M_b^\top h} = e^{-\lambda h} \sum_{k < m} \frac{h^k}{k!} (N^\top)^k \otimes R(\omega h)^\top$, which is the stated polynomial-times-sinusoid form. A concrete instance: $C = -2\mathbf{1} + \omega J$ with $\omega = 1$, $m = 2$, $V_\infty = \mathbf{1}_4$ and $\Sigma \Sigma^\top = -(M^\top + M) \succ 0$ satisfies every hypothesis and has entries $\pm h e^{-2h} \sin h$, which are not pure sinusoids. $\square$

At the reference calibration $(\kappa, \theta^*) = (4, 30)$, the required and proposed decay rates already cross between modes 2 and 3, after which the mismatch grows quadratically with mode number. With the essentially-affine caveat of Remark 2.10, this gives a direct restriction on the joint maturity and time spectra. The cylindrical model violates it maximally: its maturity spectrum has constant real part and nonzero imaginary parts; its time spectrum is real with n^2 -growing magnitude (Figure 1b, Figure 2a).




Figure 1: (a) The model's first three eigenmodes: damped oscillatory quasi-exponentials with envelope $e^{-\theta/\kappa}$. (b) The two spectra: maturity-domain $-1/\kappa \pm i\omega_n$ (vertical line) versus the model's time-domain OU rates $-\lambda_n$ (real axis); arrows mark the pairing no-arbitrage would identify — the time spectrum is the image of the maturity spectrum under the symbol p , collapsed to the real axis because the model's spectral line is the symmetry axis of the parabola. (c) The mismatch ratio $\kappa\lambda_n$ against the arbitrage-free requirement (dashed).

Remark 4.4 is a statement about factor coordinates. The remainder of the section turns it into two statements one can act on. At every truncation the defect is exposed by a finite panel of traded bonds (§4.4), and its size does not wash out under refinement (§4.5). §4.6 then asks what survives in the infinite model, where the sine system is complete and no coordinate is silent: an energy obstruction does. §4.7 closes by giving the silent coordinates noise of their own, which turns the unreachable drift into a price.

4.4 A finite-bond witness

The range defect is not merely a coordinate-level failure. At every finite truncation it can be separated by a finite panel of traded bonds. For the statement, put

$$b_n^s(\tau) := \sqrt{\frac{2}{\theta^*}} e^{-\tau/\kappa} \sin(\omega_n \tau), \quad b_n^c(\tau) := \sqrt{\frac{2}{\theta^*}} e^{-\tau/\kappa} \cos(\omega_n \tau),$$

and define the integrated sine loading and the integrated state-drift defect by

$$F_n(\tau) := \int_0^\tau b_n^s(u) du, \quad H_n(\tau) := b_n^s(\tau) + \lambda_n F_n(\tau), \quad (10)$$

and write $q^N(\tau) := \sigma_0^2 \sum_{n \leq N} b_n^s(\tau) F_n(\tau)$ for the truncation's own convexity, so that $\int_0^\tau q^N = \frac{1}{2} \sigma_0^2 \sum_n F_n(\tau)^2$. (This is $\mathbf{B}^\top C \mathbf{I}$ of Proposition 2.7 evaluated on the sine-only data, and agrees with the completion's convexity in Definition 5.1.)

With N Brownian shocks, $N + 1$ bonds generically leave one instantaneous zero-volatility portfolio. In the inconsistent truncation that portfolio retains nonzero discounted drift. The proof chooses its sign from the observed state and stops before the sign can reverse, turning the rank defect into an admissible arbitrage.

Theorem 4.6 (A finite-bond witness of the obstruction). *Fix $N \geq 1$ and consider the actual N -sine truncation*

$$r_t(\tau) = \varphi_t(\tau) + \sum_{n=1}^N b_n^s(\tau) x_{n,t}, \quad dx_{n,t} = -\lambda_n x_{n,t} dt + \sigma_0 dW_t^n, \quad (11)$$

with deterministic initial state and a deterministic, jointly C^1 sector $\varphi(t, \tau)$. Let $B_t = \exp(\int_0^t r_s(0) ds)$ and $\tilde{P}(t, T) = P(t, T)/B_t$.

For every nonempty open interval $J \subset (0, \theta^*)$, define the ordered simplex

$$J_{\uparrow}^{N+1} := \{(\tau_1, \dots, \tau_{N+1}) \in J^{N+1} : \tau_1 < \dots < \tau_{N+1}\}.$$

An open dense subset of full Lebesgue measure in $J_{\uparrow}^{N+1}$ has the following property. At every deterministic $t_0 > 0$, the fixed-calendar-maturity bonds $T_j = t_0 + \tau_j$, together with the savings account, support an admissible self-financing strategy with zero initial wealth, nonnegative wealth at all times, and strictly positive terminal wealth almost surely. Thus the finite bond market violates NA (hence NFLVR) and admits no equivalent local martingale measure.

The count is generically minimal for this instantaneous zero-volatility construction: $N + 1$ denotes $N + 1$ risky bonds, in addition to the savings account.

Proof. Write

$$a_0(t, \tau) := \int_0^\tau (\dot{\varphi}_t(u) - \varphi'_t(u) - q^N(u)) du.$$

With F_n, H_n as in (10), the state-linear difference between the model's drift and HJM transport in (11) is

$$\sum_{n=1}^N \left[\left(\frac{1}{\kappa} - \lambda_n \right) b_n^s(\tau) - \omega_n b_n^c(\tau) \right] x_{n,t}.$$

Since $(b_n^s)' = -\kappa^{-1} b_n^s + \omega_n b_n^c$, its maturity integral is $-H_n$. The bond calculation in Lemma 2.1, including $\int_0^\tau q^N = \frac{1}{2} \sigma_0^2 \sum_n F_n(\tau)^2$, therefore gives

$$\frac{d\tilde{P}(t, T)}{\tilde{P}(t, T)} = \left[-a_0(t, T-t) + \sum_{n=1}^N H_n(T-t) x_{n,t} \right] dt - \sigma_0 \sum_{n=1}^N F_n(T-t) dW_t^n. \quad (12)$$

The deterministic HJM defect enters as the intercept a_0 .

Fix any $k \leq N$. The analytic functions $F_1, \dots, F_N, H_k$ are linearly independent on every open J . Indeed, differentiation first proves independence of the F_n from that of the distinct sine modes. If H_k lay in their span, differentiating would put

$$H'_k = \left(\lambda_k - \frac{1}{\kappa} \right) b_k^s + \omega_k b_k^c$$

in the finite sine span, impossible because $\omega_k > 0$ and the sine–cosine characters are linearly independent on every open interval. The elementary evaluation lemma for linearly independent functions now supplies distinct $\tau_1, \dots, \tau_{N+1} \in J$ such that

$$\det \begin{bmatrix} F_1(\tau_1) & \cdots & F_N(\tau_1) & H_k(\tau_1) \\ \vdots & & \vdots & \vdots \\ F_1(\tau_{N+1}) & \cdots & F_N(\tau_{N+1}) & H_k(\tau_{N+1}) \end{bmatrix} \neq 0. \quad (13)$$

For completeness, the evaluation lemma follows by induction: if no final evaluation point made the determinant nonzero, the last function would be a fixed linear combination of the preceding ones. The determinant in (13) is a nonzero real-analytic function of the maturity tuple, so its

nonzero set is open and dense in the ordered simplex. Since the zero set of a nonvanishing real-analytic function is Lebesgue-null, that set also has full measure, as the theorem asserts.

Set $T_j = t_0 + \tau_j$ and, for t near t_0 , define

$$\mathbf{F}_{jn}(t) := F_n(T_j - t), \quad \ell_j(t) := (-1)^{j+N+1} \det \mathbf{F}^{(j)}(t),$$

where $\mathbf{F}^{(j)}$ deletes row j . Laplace expansion gives $\ell(t)^\top \mathbf{F}(t) = 0$. By (13) and continuity, there is a deterministic $\varepsilon > 0$ with $t_0 + \varepsilon < \min_j T_j$ such that $\text{rank } \mathbf{F}(t) = N$ and

$$\ell(t)^\top (H_k(T_1 - t), \dots, H_k(T_{N+1} - t))^\top \neq 0 \quad (14)$$

throughout $[t_0, t_0 + \varepsilon]$.

Let $\mu(t)$ be the vector of discounted relative drifts in (12) and put $g_t := \ell(t)^\top \mu(t)$. By (14), the coefficient of x_{k,t_0} in g_{t_0} is nonzero. Since

$$\text{Var}(x_{n,t_0}) = \frac{\sigma_0^2}{2\lambda_n} (1 - e^{-2\lambda_n t_0}) > 0,$$

g_{t_0} is a nondegenerate affine Gaussian, even in the presence of the deterministic intercept a_0 , and hence $\mathbb{P}(g_{t_0} = 0) = 0$.

Set $\zeta := \text{sgn}(g_{t_0})$ and

$$\varrho := \inf \{t \geq t_0 : \zeta g_t \leq \tfrac{1}{2}|g_{t_0}|\} \wedge (t_0 + \varepsilon).$$

Here $g_t = \ell(t)^\top \mu(t)$ is jointly continuous in (t, ω) : ℓ and a_0 are deterministic and continuous, and μ is affine in the continuous state $\mathbf{X}_t$ with continuous deterministic coefficients. Since $\zeta g_{t_0} = |g_{t_0}| > 0$ a.s., continuity of $t \mapsto \zeta g_t$ gives $\varrho > t_0$ almost surely. Use the predictable discounted dollar exposures

$$h_j(t) := \zeta \ell_j(t) \mathbf{1}_{(t_0, \varrho]}(t),$$

equivalently the bond units

$$\vartheta_j(t) := \frac{\zeta \ell_j(t)}{\tilde{P}(t, T_j)} \mathbf{1}_{(t_0, \varrho]}(t).$$

Here ζ is $\mathcal{F}_{t_0}$ -measurable, and the continuous hitting time ϱ is predictable, so the exposures are predictable. Because every $\tilde{P}(\cdot, T_j)$ is strictly positive and continuous and ℓ is deterministic and continuous, the bond units are locally bounded and hence semimartingale-integrable. Their discounted gain, started from zero, satisfies exactly

$$dG_t = \sum_{j=1}^{N+1} \vartheta_j(t) d\tilde{P}(t, T_j) = \zeta g_t \mathbf{1}_{(t_0, \varrho]}(t) dt,$$

because $\ell^\top \mathbf{F} = 0$. It has no Brownian term and

$$G_t \geq 0, \quad G_{t_0+\varepsilon} = G_\varrho \geq \tfrac{1}{2}|g_{t_0}|(\varrho - t_0) > 0 \quad \text{a.s.}$$

The risky positions are financed through the discounted cash holding

$$\vartheta_t^0 := G_t - \sum_j \vartheta_j(t) \tilde{P}(t, T_j).$$

After ϱ the positive wealth is left in cash. This defines a self-financing admissible arbitrage. Finally, for $m \leq N$ generic bonds the $m \times N$ evaluation matrix of the independent F_n has row rank m , hence no nonzero left-null portfolio. This proves the generic-minimality claim. $\square$

Remark 4.7 (What the construction needs: a recipe, not a computation). The proof used three properties of the model and nothing else, so the route generalises. Let a factor term-structure model have (i) loadings real-analytic on some open maturity interval J ; (ii) a state-linear discrepancy $\sum_n c_n(\tau)x_{n,t}$ between its drift and HJM transport, whose maturity integral $\int_0^\cdot c_k$ fails to lie in the span of the integrated volatility loadings $\{\int_0^\cdot \sigma_n\}_{n \leq N}$ as functions on J — equivalently, a discrepancy no Girsanov tilt can reach; and (iii) a state making the *surviving drift functional* $g_{t_0} = \ell(t_0)^\top \mu(t_0)$ nondegenerate, in the sense that its law is atom-free at 0. Nondegeneracy of the individual coordinates is not enough: if the cofactor pairing leaves the combination $a = (1, -1)$ while $x_{1,t_0} = x_{2,t_0}$ almost surely, both coordinates are nondegenerate and $a^\top \mathbf{X}_{t_0}$ vanishes identically. A state law with positive-definite covariance on the relevant coordinates — as in Theorem 4.6, where the x_n are independent with positive variance — is a convenient sufficient condition. Then the evaluation-determinant argument supplies an open dense set of maturity tuples, the cofactor portfolio annihilates the volatility matrix identically, and the sign-and-stopping argument converts the surviving drift into an admissible arbitrage in $N + 1$ bonds. Nothing in this uses the sine basis, the cylindrical noise, or the particular symbol. The passage from “the drift condition fails in a direction outside range Σ ” to an explicit finite portfolio is therefore available for any such model. And it is constructive, where the fundamental theorem of asset pricing would give only existence.

4.5 The size of the obstruction

The next result measures the size of the unreachable drift. Its quotient class is intrinsic; the displayed scalar uses the model’s own normalization of the coordinates, and Remark 4.9 says what it does and does not measure.

Proposition 4.8 (Normalized-coordinate size of the unspanned drift). *Order the n th envelope state as (s_n, c_n) , with the sine coordinate first, and let*

$$K_N^C := \bigoplus_{n=1}^N (-\lambda_n \mathbf{1}_2), \quad M_N := \bigoplus_{n=1}^N M_n, \quad \Sigma_N := \sigma_0(\mathbf{e}_1^s, \dots, \mathbf{e}_N^s).$$

The invariant object is the quotient class

$$\mathfrak{o}_N(x) := [(K_N^C - M_N^\top)x] \in \mathbb{R}^{2N} / \text{range } \Sigma_N.$$

In the canonically normalized sine-cosine coordinates above, equip the quotient with the induced Euclidean norm and put $d_n := \kappa^{-1} - \lambda_n$. Then

$$\mathfrak{d}_N(x)^2 := \inf_{\gamma \in \mathbb{R}^N} \left\| (K_N^C - M_N^\top)x - \Sigma_N \gamma \right\|_2^2 = \sum_{n=1}^N (-\omega_n s_n + d_n c_n)^2. \quad (15)$$

For the stationary embedded sine truncation, $c_n = 0$ and the s_n are independent centred Gaussians with variance $\sigma_0^2/(2\lambda_n)$. Thus, for independent standard normal variables Z_n ,

$$\mathfrak{d}_N(X)^2 \stackrel{d}{=} \sum_{n=1}^N \frac{\sigma_0^2 \kappa \omega_n^2}{1 + \kappa^2 \omega_n^2} Z_n^2, \quad (16)$$

and consequently, from (16),

$$\mathbb{E} \mathfrak{d}_N(X)^2 = \frac{\sigma_0^2}{\kappa} \left[N - \sum_{n=1}^N \frac{1}{1 + \kappa^2 \omega_n^2} \right] \quad (17)$$

$$= \frac{\sigma_0^2}{\kappa} N - \frac{\sigma_0^2}{2\kappa} \left[\frac{\theta^*}{\kappa} \coth\left(\frac{\theta^*}{\kappa}\right) - 1 \right] + O(N^{-1}). \quad (18)$$

Under the natural nested coupling of the stationary coordinates,

$$\frac{\mathfrak{d}_N(X)^2}{N} \longrightarrow \frac{\sigma_0^2}{\kappa} \quad \text{almost surely and in } L^2. \quad (19)$$

Thus the normalized squared obstruction is extensive: it approaches σ_0^2/κ per retained mode rather than disappearing with refinement.

Proof. On block n ,

$$K_n^C - M_n^\top = \begin{pmatrix} d_n & \omega_n \\ -\omega_n & d_n \end{pmatrix}.$$

The volatility range is exactly the sine coordinate in every block. Minimization over γ_n removes the first row and leaves $-\omega_n s_n + d_n c_n$, proving (15). In stationarity,

$$\mathbb{E}(\omega_n^2 s_n^2) = \frac{\sigma_0^2 \omega_n^2}{2\lambda_n} = \frac{\sigma_0^2}{\kappa} \left(1 - \frac{1}{1 + \kappa^2 \omega_n^2} \right),$$

which gives (17). The identity

$$\sum_{n=1}^{\infty} \frac{1}{1 + (\kappa \pi n / \theta^*)^2} = \frac{1}{2} \left[\frac{\theta^*}{\kappa} \coth\left(\frac{\theta^*}{\kappa}\right) - 1 \right]$$

and its $O(N^{-1})$ tail give (18). Finally the independent summands $\omega_n^2 s_n^2$ have uniformly bounded variances and means converging to σ_0^2/κ . Kolmogorov's strong law gives the almost-sure limit in (19); the same variance bound gives L^2 convergence. $\square$

Remark 4.9 (What the modulus does and does not measure). The quotient class $\mathfrak{o}_N$ and its nonvanishing are invariant under invertible factor changes; the Euclidean number $\mathfrak{d}_N$ is tied to the normalized sine-cosine coordinates and changes under a non-orthogonal rescaling. It quantifies only the state-linear drift component outside the Girsanov range, not the deterministic convexity defect. It is therefore neither a market price of risk, a Sharpe ratio, a hedging loss, nor a monetary reserve. A security-output norm would additionally require a declared maturity panel and portfolio metric; Theorem 4.6 supplies the invariant fact needed here, namely that a generic finite panel separates the quotient class.

That completes the finite- N diagnosis. What survives it in the infinite model is the next question.

4.6 The infinite model: an energy obstruction

In the infinite cylindrical model the sine system is complete in the state space, so no literal “noiseless coordinate” survives the limit and Remark 4.4 does not upgrade termwise. What survives is an energy obstruction.

Proposition 4.10 (The infinite model: no square-integrable Girsanov repair). *(i) The stationary field’s maturity increments obey $\mathbb{E}(X_t(\theta + h) - X_t(\theta))^2 = \frac{\sigma_0^2}{\kappa} e^{-2\theta/\kappa} |h| + O(h^2)$ at every interior θ . The field is Brownian-rough in maturity, hence a.s. nowhere differentiable in θ , and $X_t \notin C^1$ a.s. (ii) On the completed natural filtration of the driving cylindrical Brownian motion, every equivalent change of measure is represented by a predictable Girsanov density γ satisfying*

$$\int_0^T \|\gamma_t\|_H^2 dt < \infty \quad \text{a.s.},$$

together with the condition that its stochastic exponential is a true martingale. Thus the relevant integrand space is the pathwise space $L^2([0, T]; H)$, not $L^2([0, T]; H_w)$. For the stationary solution, and also for every deterministic $X_0 \in H$ independent of the future noise, at each $t > 0$

$$\sum_{n \geq 1} \omega_n^4 |x_n(t)|^2 = \infty \quad \text{a.s.} \quad (20)$$

Consequently $\partial_\theta^2 X_t \notin H$ a.s. for every $t > 0$, and with probability one this holds for Lebesgue-a.e. $t \in (0, T)$. Hence there is no predictable $\gamma \in L^2([0, T]; H)$ for which

$$\sigma_0 \gamma_t = (\tilde{A} - \partial_\theta) X_t - q_t = \frac{\kappa}{2} \partial_\theta^2 X_t - q_t$$

holds distributionally for a.e. t , whenever the proposed convexity term q_t is H -valued. In particular, an equivalent measure change cannot implement this repair even if the drift identity is required only against the test core $C_c^\infty(0, \theta^*)$. The spatial assertion in (i) is the stationary stochastic-heat regularity; see [41, 33]. Part (ii) is proved below.

Proof. The map $U : H \rightarrow L^2(0, \theta^*)$, $(Uf)(\theta) = e^{\theta/\kappa} f(\theta)$, is unitary and $Ue_n = \phi_n := \sqrt{2/\theta^*} \sin(\omega_n \theta)$. In stationarity, the coefficients of $Y := UX$ have variances

$$\frac{\sigma_0^2}{2\lambda_n} = \frac{\sigma_0^2}{\kappa(\omega_n^2 + \kappa^{-2})}.$$

Therefore

$$\text{Cov}(Y(\theta), Y(u)) = \frac{\sigma_0^2}{\kappa} G_{-\partial_\theta^2 + \kappa^{-2}}^D(\theta, u),$$

where G^D is the Dirichlet Green kernel. Its first derivative has the unit diagonal jump, so the covariance increment identity gives $\mathbb{E}[(Y(\theta + h) - Y(\theta))^2] = (\sigma_0^2/\kappa)|h| + O(h^2)$. Since $X(\theta) = e^{-\theta/\kappa} Y(\theta)$, part (i)’s displayed asymptotic follows. The Gaussian process Y is the massive Brownian (Ornstein–Uhlenbeck) bridge in the spatial coordinate; its quadratic variation is $(\sigma_0^2/\kappa)d\theta$. Thus it, and its product with the positive smooth envelope $e^{-\theta/\kappa}$, is a.s. nowhere differentiable.

For (ii), moreover,

$$UAU^{-1} = \frac{\kappa}{2} \partial_\theta^2 - \frac{1}{2\kappa}$$

with Dirichlet domain. For deterministic initial data, the independent Gaussian noise parts of $x_n(t)$ have variances

$$v_n(t) = \frac{\sigma_0^2}{2\lambda_n} (1 - e^{-2\lambda_n t}) \asymp \omega_n^{-2} \quad (t > 0, n \rightarrow \infty);$$

in stationarity, $v_n = \sigma_0^2/(2\lambda_n)$ exactly. Writing $x_n(t) = m_n(t) + \sqrt{v_n(t)}Z_n$ with independent standard normals, the events $|x_n(t)| \geq \frac{1}{2}\sqrt{v_n(t)}$ have probabilities bounded below uniformly in n (the centred interval has maximal Gaussian mass when $m_n(t) = 0$). Independence and the second Borel–Cantelli lemma therefore imply that infinitely many occur. Their terms in $\sum_n \omega_n^4 x_n(t)^2$ are bounded below by a positive multiple of ω_n^2 , proving (20).

Zero traces are needed before the sine criterion may be applied, and they do not follow from (20) alone: $f \equiv 1$ has $f'' = 0 \in H$ while its sine coefficients are of order n^{-1} , so its weighted sum diverges. They do hold here. In stationarity part (i) exhibits $Y = UX_t$ as the massive Brownian bridge on $[0, \theta^*]$, which has a continuous modification vanishing at both endpoints; for deterministic $X_0 \in H$, write $X_t = e^{tA}X_0 + Z_t$ where $Z_t := X_t^{\text{stat}} - e^{tA}X_0^{\text{stat}}$ is a version of the stochastic convolution. The semigroup terms lie in $D(A) \subset H^2 \cap H_0^1$ for $t > 0$ (A is self-adjoint and negative after conjugation, hence analytic), and the bridge terms vanish at the endpoints, so X_t is continuous with $X_t(0) = X_t(\theta^*) = 0$ a.s.

Now suppose $\partial_\theta^2 X_t \in H$. One-dimensional elliptic regularity then gives $X_t \in H^2$, and with the zero traces just established $UX_t \in H^2 \cap H_0^1$ (multiplication by $e^{\pm\theta/\kappa}$ preserves H^2 on a finite interval). The Dirichlet sine characterisation — $UX_t \in H^2 \cap H_0^1$ if and only if $\sum_n (1 + \omega_n^4)|x_n(t)|^2 < \infty$ — then contradicts (20). Hence $\partial_\theta^2 X_t \notin H$. Since this holds with probability one for each $t > 0$, Tonelli's theorem applied to the indicator of the exceptional H^2 event gives the simultaneous statement for Lebesgue-a.e. t , almost surely. If the displayed weak drift identity held with $\gamma_t, q_t \in H$, its right-hand side would give an H representative of $\partial_\theta^2 X_t$ as a distribution, contradicting the preceding conclusion. Finally, on the completed natural Brownian filtration, Girsanov's theorem requires $\int_0^T \|\gamma_t\|_H^2 dt < \infty$ almost surely, in addition to the true-martingale condition for the density exponential, so such a measure change is unavailable. $\square$

That is the obstruction under the model's own noise. One question remains before the diagnosis is complete, and the next subsection settles it: what changes if the silent coordinates are given noise of their own, so that the defect becomes reachable by a change of measure.

4.7 When the silent coordinates are loaded: a market-price-of-risk cost

Remark 4.4 and Proposition 4.10 rest on the same fact: under the model's sine-only loading the defect leaves the range of the volatility, so no tilt can reach it. Suppose instead that every block is loaded isotropically, with intensities $\sigma_n > 0$ and $\sum_n \omega_n^2 \sigma_n^2 < \infty$, and use the damping-compatible weight (1). This loading adds one Brownian source per block: it is a second completion, compared with the sine-only one in Remark 5.3(iii). With it the defect lies inside the volatility range, Girsanov's theorem applies at every finite order, and the question becomes what the tilt costs and whether the cost stays finite on the infinite envelope. The answer is a dichotomy governed by one number.

For an entire function p with $p(\bar{z}) = \overline{p(z)}$ the induced drift on block n is $p(M_n)^\top$: the proof of Proposition 4.2(a) applies to the power series of p , which converges on the two-dimensional block.

Since $M_n = -\kappa^{-1}\mathbf{1} + \omega_n J$ with $J^2 = -\mathbf{1}$, every entire f satisfies $f(M_n) = \operatorname{Re} f(z_n) \mathbf{1} + \operatorname{Im} f(z_n) J$ with $z_n := -\kappa^{-1} + i\omega_n$. Write

$$\delta_n := p(z_n) - z_n, \quad k_n := -\operatorname{Re} p(z_n), \quad \mathfrak{e}_n := \frac{|\delta_n|^2}{k_n}, \quad \mathcal{E} := \sum_{n \geq 1} \mathfrak{e}_n, \quad (21)$$

and let $\Xi_n := p(M_n)^\top - M_n^\top$ be the *defect block*, the difference between the induced drift and the arbitrage-free drift of Proposition 2.7. Then $\Xi_n^\top \Xi_n = |\delta_n|^2 \mathbf{1}_2$, and the symmetric part of $p(M_n)^\top$ is $-k_n \mathbf{1}_2$: k_n is the mean-reversion rate of block n under the induced dynamics and κ^{-1} its rate under the arbitrage-free dynamics. We call $\mathfrak{e}_n$ the *defect energy* of block n and $\mathcal{E}$ the defect energy of the envelope. For the model's symbol $p(z) = z + \frac{\kappa}{2} z^2$ one has $p(z_n) = -\lambda_n$ (Proposition 3.1), hence $\delta_n = -(\lambda_n + z_n)$, $|\delta_n|^2 = d_n^2 + \omega_n^2 = \lambda_n^2$ with $d_n = \kappa^{-1} - \lambda_n$ as in Proposition 4.8, $k_n = \lambda_n$, $\mathfrak{e}_n = \lambda_n$, and Ξ_n is the block displayed in Remark 4.4.

All laws below live on the canonical product path space $\Omega := \prod_{n \geq 1} C([0, T]; \mathbb{R}^2)$, $0 < T < \infty$, with its product σ -algebra $\mathcal{F}_T$. The coordinate process is $x = (x_n)_{n \geq 1}$ with $x_n = (s_n, c_n)$; its raw canonical filtration is augmented under each law separately where needed, and $\mathcal{F}_T^{(N)} := \sigma\{x_n(s) : n \leq N, s \leq T\}$ is the σ -field of the first N blocks. Writing b_n^s, b_n^c for the two components of $\mathbf{B}_n$, direct differentiation and (1) give

$$\|b_n^s\|_w^2 + \|b_n^c\|_w^2 = \frac{2}{\theta^*} [1 + (\omega_n^2 + \kappa^{-2}) C_w].$$

Consequently $\sum_n \omega_n^2 \sigma_n^2 < \infty$ makes $\sum_n \mathbf{B}_n^\top x_n$ an H_w -valued process under the arbitrage-free dynamics. The same is true under the induced dynamics whenever $\inf_n k_n > 0$, since each coordinate variance is then bounded by a constant multiple of σ_n^2 . Nothing in the proposition itself requires the coordinate sequence to define a curve: its statements concern the two product laws.

Proposition 4.11 (Girsanov energy and the Kakutani dichotomy). *Under the isotropic loading above, let $\mathbb{P}^{\text{def}}$ and $\mathbb{P}^{\text{NA}}$ be the laws on $\mathcal{F}_T$, $0 < T < \infty$, of the factor process started at the origin, under the induced drift $p(M)^\top$ and under the arbitrage-free drift $M^\top$ respectively, with the same noise. Assume $k_n > 0$ for all n , except where (c) says otherwise; in (a) this assumption enters only the stationary statement.*

(a) (Finite order.) *Let $P_N := \mathbb{P}^{\text{def}}|_{\mathcal{F}_T^{(N)}}$ and $Q_N := \mathbb{P}^{\text{NA}}|_{\mathcal{F}_T^{(N)}}$. For every N , $P_N \sim Q_N$ with*

$$\frac{dQ_N}{dP_N} = \exp\left(-\sum_{n \leq N} \int_0^T \gamma_{n,t} \cdot d\mathbf{W}_t^n - \frac{1}{2} \sum_{n \leq N} \int_0^T |\gamma_{n,t}|^2 dt\right), \quad \gamma_{n,t} = \sigma_n^{-1} \Xi_n \begin{pmatrix} s_{n,t} \\ c_{n,t} \end{pmatrix},$$

where $\mathbf{W}^n$ is the $\mathbb{P}^{\text{def}}$ -Brownian motion of block n . In $\mathbb{P}^{\text{def}}$ -stationarity $\mathbb{E} \sum_{n \leq N} |\gamma_{n,t}|^2 = \sum_{n \leq N} \mathfrak{e}_n$, whatever the intensities. The Hellinger affinity of P_N and Q_N is $\prod_{n \leq N} \mathfrak{h}_n$ with

$$\mathfrak{h}_n = \frac{e^{\bar{k}_n T}}{\cosh(\mathfrak{m}_n T) + (\bar{k}_n / \mathfrak{m}_n) \sinh(\mathfrak{m}_n T)}, \quad \bar{k}_n := \frac{1}{2}(k_n + \kappa^{-1}), \quad \mathfrak{m}_n^2 := \bar{k}_n^2 + \frac{1}{4}|\delta_n|^2. \quad (22)$$

(b) (Infinite envelope.) On $\mathcal{F}_T$,

$$\mathbb{P}^{\text{def}} \sim \mathbb{P}^{\text{NA}} \iff \mathcal{E} < \infty \iff \sum_n |\delta_n|^2 < \infty \iff \Xi := \bigoplus_n \Xi_n \text{ is Hilbert-Schmidt,}$$

Ξ acting on $\ell^2(\mathbb{N}; \mathbb{R}^2)$; otherwise the two laws are mutually singular, with $\int_0^T |\gamma_t|^2 dt = \infty$ $\mathbb{P}^{\text{def}}$ -a.s., where $|\gamma_t|^2 := \sum_{n \geq 1} |\gamma_{n,t}|^2$. In the singular case no measure equivalent to $\mathbb{P}^{\text{def}}$ gives the factor process the dynamics of Proposition 2.7.

(c) Equivalence forces $\delta_n \rightarrow 0$, hence $k_n \rightarrow \kappa^{-1}$. It fails for every polynomial symbol of degree at least two. If $k_n \leq 0$ on infinitely many blocks, the two product laws are still defined on Ω and are mutually singular there, although the coordinate sequence need not then define an H_w -valued curve.

(d) (The model.) For $p(z) = z + \frac{\kappa}{2}z^2$,

$$\mathcal{E}_N := \sum_{n \leq N} \mathfrak{e}_n = \sum_{n \leq N} \lambda_n = \frac{N}{2\kappa} + \frac{\kappa\pi^2 N(N+1)(2N+1)}{12\theta^{*2}},$$

and, with $c_1 := \frac{1}{\sqrt{2}} - \frac{1}{2}$,

$$e^{-T\mathcal{E}_N/4} \leq \prod_{n \leq N} \mathfrak{h}_n \leq 2^N e^{(1-1/\sqrt{2})NT/(2\kappa)} e^{-c_1 T\mathcal{E}_N},$$

so that $-\log \prod_{n \leq N} \mathfrak{h}_n \sim c_1 T \mathcal{E}_N \sim c_1 \kappa \pi^2 T N^3 / (6\theta^{*2})$. The tilted truncations are therefore equivalent at every order and singular in the limit, and their affinity decays like e^{-cN^3} on both sides.

Part (d) is the quantitative counterpart of Remark 4.4. With the sine-only loading the two laws are inequivalent at every order because the defect is unreachable; with the isotropic loading they are equivalent at every order, at a cost growing like N^3 , and inequivalent in the limit. The two loadings agree on the infinite verdict and differ in what they say about truncations: an impossibility against a price.

Proof. (a) On the first N blocks the drift is linear in the state, so the stochastic exponential of the linear $(\gamma_n)_{n \leq N}$ is a true martingale (Beneš' criterion, [31, Cor. 3.5.16]). Writing the factor equation of those blocks as $dx = p(M)^\top x dt + \Sigma d\mathbf{W} = M^\top x dt + \Sigma (d\mathbf{W} + \gamma dt)$ identifies $\mathbf{W} + \int \gamma dt$ as a $2N$ -dimensional Brownian motion under the tilted measure on $\mathcal{F}_T^{(N)}$, whose density is the displayed exponential; neither this nor the affinity formula below uses $k_n > 0$. Block n of $p(M)^\top$ is normal with symmetric part $-k_n \mathbf{1}$, so its stationary covariance solves the Lyapunov equation and equals $\sigma_n^2 / (2k_n) \mathbf{1}_2$; hence $\mathbb{E}|\gamma_n|^2 = \sigma_n^{-2} \text{tr}(\Xi_n^\top \Xi_n) \sigma_n^2 / (2k_n) = |\delta_n|^2 / k_n = \mathfrak{e}_n$. For the affinity let $b^{\text{def}}, b^{\text{NA}}$ be the two drifts and $\mathbb{P}^{\text{mid}}$ the law with the midpoint drift $\frac{1}{2}(b^{\text{def}} + b^{\text{NA}})$. Since $\frac{1}{2}(|b^{\text{def}}|^2 + |b^{\text{NA}}|^2) - |\frac{1}{2}(b^{\text{def}} + b^{\text{NA}})|^2 = \frac{1}{4}|b^{\text{def}} - b^{\text{NA}}|^2$, the Girsanov densities against the driftless reference law give $\sqrt{d\mathbb{P}^{\text{def}} d\mathbb{P}^{\text{NA}}} = d\mathbb{P}^{\text{mid}} \exp(-\frac{1}{8} \int_0^T |\Sigma^{-1} \Xi x_t|^2 dt)$. Block n of the midpoint drift is normal with symmetric part $-\bar{k}_n \mathbf{1}$; its skew part rotates (s_n, c_n) without changing the modulus, so $\mathfrak{h}_n = \mathbb{E}^{\text{mid}} \exp(-\frac{1}{8} |\delta_n|^2 \int_0^T |y_t|^2 dt)$ with y a two-dimensional Ornstein–Uhlenbeck process of rate $\bar{k}_n$ and unit intensity started at 0 (the intensity σ_n cancels). For a one-dimensional such process the Riccati equation for the exponent gives

$\mathbb{E} \exp(-\vartheta \int_0^T y_t^2 dt) = e^{\bar{k}T/2} (\cosh \mathbf{m}T + (\bar{k}/\mathbf{m}) \sinh \mathbf{m}T)^{-1/2}$ with $\mathbf{m}^2 = \bar{k}^2 + 2\vartheta$; squaring for the two independent coordinates and putting $\vartheta = |\delta_n|^2/8$ gives (22).

(b) Under both laws the blocks are independent, so on $\mathcal{F}_T$ both laws are product measures over n with equivalent factors by (a), and Kakutani's theorem [30] states that they are equivalent if $\prod_n \mathfrak{h}_n > 0$ and mutually singular otherwise. *Sufficiency.* By Jensen, $\mathfrak{h}_n = \mathbb{E}^{\text{def}} \sqrt{d\mathbb{P}_n^{\text{NA}}/d\mathbb{P}_n^{\text{def}}} \geq \exp(-\frac{1}{2}\text{KL}_n)$ with $\text{KL}_n = \frac{1}{2}\mathbb{E}^{\text{def}} \int_0^T |\gamma_{n,t}|^2 dt$; from the origin $\mathbb{E}^{\text{def}} |x_{n,t}|^2 = (\sigma_n^2/k_n)(1 - e^{-2k_nt})$, so $\text{KL}_n = \frac{1}{2}\mathfrak{e}_n [T - (1 - e^{-2k_nT})/(2k_n)] \leq \frac{1}{2}\mathfrak{e}_n T$ and $1 - \mathfrak{h}_n \leq \frac{1}{4}\mathfrak{e}_n T$. Thus $\mathcal{E} < \infty$ gives $\sum_n (1 - \mathfrak{h}_n) < \infty$ and equivalence. *Necessity.* By Kakutani's criterion it suffices to show that $\mathcal{E} = \infty$ forces $\sum_n (1 - \mathfrak{h}_n) = \infty$, and this is done block by block. Put $Y_n := \int_0^T |\gamma_{n,t}|^2 dt$. The midpoint representation in the proof of (a) reads $\mathfrak{h}_n = \mathbb{E}^{\text{mid}} e^{-Y_n/8}$, so $1 - \mathfrak{h}_n = \mathbb{E}^{\text{mid}} (1 - e^{-Y_n/8}) \geq (1 - e^{-1}) \mathbb{E}^{\text{mid}} \min(Y_n/8, 1)$. Under $\mathbb{P}^{\text{mid}}$, Y_n is a positive quadratic functional of a centred Gaussian process, hence by the Karhunen–Loève expansion a sum $\sum_i a_i Z_i^2$ with $a_i \geq 0$ and independent standard normals Z_i ; so $\mathbb{E} Y_n^2 \leq 3(\mathbb{E} Y_n)^2$ and Paley–Zygmund gives $\mathbb{P}^{\text{mid}}(Y_n \geq \frac{1}{2}\mathbb{E}^{\text{mid}} Y_n) \geq \frac{1}{12}$, the step already used in Proposition 3.3(c); hence $\mathbb{E}^{\text{mid}} \min(Y_n/8, 1) \geq \frac{1}{12} \min(\frac{1}{16}\mathbb{E}^{\text{mid}} Y_n, 1)$. Block n under $\mathbb{P}^{\text{mid}}$ is an Ornstein–Uhlenbeck process with symmetric part $-\bar{k}_n \mathbf{1}$ started at the origin, so $\mathbb{E}^{\text{mid}} |x_{n,t}|^2 = \sigma_n^2(1 - e^{-2\bar{k}_nt})/\bar{k}_n$ and

$$\mathbb{E}^{\text{mid}} Y_n = \frac{|\delta_n|^2}{\bar{k}_n} T \varphi(\bar{k}_n T), \quad \varphi(u) := 1 - \frac{1 - e^{-2u}}{2u} \in (0, 1), \quad \varphi(u) \geq \frac{1}{3} \min(u, 1),$$

φ increasing. Split the index set at $k_n \geq \frac{1}{2\kappa}$. If $\sum \mathfrak{e}_n = \infty$ over the blocks with $k_n \geq \frac{1}{2\kappa}$, then there $\bar{k}_n \geq \frac{3}{4\kappa}$ and $\frac{2k_n}{k_n + \kappa^{-1}} \geq \frac{2}{3}$, so $\mathbb{E}^{\text{mid}} Y_n = \frac{2k_n}{k_n + \kappa^{-1}} \mathfrak{e}_n T \varphi(\bar{k}_n T) \geq c \mathfrak{e}_n$ with $c := \frac{2}{9} T \min(\frac{3T}{4\kappa}, 1)$, and $\sum \min(\frac{c}{16} \mathfrak{e}_n, 1) = \infty$. Otherwise $\sum \mathfrak{e}_n = \infty$ over the blocks with $k_n < \frac{1}{2\kappa}$, which are then infinitely many; there $|\delta_n| \geq \text{Re } \delta_n = \kappa^{-1} - k_n > \frac{1}{2\kappa}$ and $\bar{k}_n \in (\frac{1}{2\kappa}, \frac{3}{4\kappa})$, so $\mathbb{E}^{\text{mid}} Y_n \geq \frac{T}{3\kappa} \varphi(\frac{T}{2\kappa}) > 0$ on infinitely many blocks. In either case $\sum_n (1 - \mathfrak{h}_n) = \infty$, and Kakutani's dichotomy gives singularity. The pathwise statement follows from the same estimates under $\mathbb{P}^{\text{def}}$: the Y_n are independent and nonnegative with $\mathbb{E}^{\text{def}} Y_n = 2 \text{KL}_n = \mathfrak{e}_n T \varphi(k_n T)$; Paley–Zygmund gives $\mathbb{E}^{\text{def}} \min(Y_n, 1) \geq \frac{1}{12} \min(\frac{1}{2}\mathbb{E}^{\text{def}} Y_n, 1)$; the same split shows $\sum_n \mathbb{E}^{\text{def}} \min(Y_n, 1) = \infty$ (in the second case $\mathfrak{e}_n \varphi(k_n T) \geq |\delta_n|^2 \min(T, k_n^{-1})/3 \geq \min(T, 2\kappa)/(12\kappa^2)$); and for independent nonnegative summands this is equivalent to $\sum_n Y_n = \infty$ a.s. (the three-series theorem with Borel–Cantelli). The remaining equivalences follow from $k_n = \kappa^{-1} - \text{Re } \delta_n \leq \kappa^{-1} + |\delta_n|$, which gives $\mathfrak{e}_n \geq |\delta_n|^2/(\kappa^{-1} + |\delta_n|)$: $\mathcal{E} < \infty$ forces $\delta_n \rightarrow 0$, then $k_n \rightarrow \kappa^{-1}$ and $\mathfrak{e}_n \asymp \kappa |\delta_n|^2$, while $\|\Xi\|_{\text{HS}}^2 = 2 \sum_n |\delta_n|^2$. For the last sentence: under a measure equivalent to $\mathbb{P}^{\text{def}}$ that gives the factor process the drift $M^\top$ with the same noise, the law of the factor process would be $\mathbb{P}^{\text{NA}}$, which is singular to $\mathbb{P}^{\text{def}}$.

(c) is contained in (b): for a polynomial p of degree at least two, $|p(z_n)|/\omega_n \rightarrow \infty$, so $|\delta_n| \geq |p(z_n)| - |z_n| \rightarrow \infty$ and $\sum_n |\delta_n|^2 = \infty$. Blocks with $k_n \leq 0$ do not affect the equivalence in (a) or the affinity formula, neither of which uses $k_n > 0$; only the stationary energy does. If only finitely many such blocks occur, discard them: their factors are equivalent by (a) and do not affect Kakutani's product, so (b) applied to the remaining blocks decides the dichotomy. When infinitely many occur, the midpoint rate $\bar{k}_n = \frac{1}{2}(k_n + \kappa^{-1})$ is at most $\frac{1}{2\kappa}$ there, so $\mathbb{E}^{\text{mid}} |x_{n,t}|^2 = 2\sigma_n^2 \int_0^t e^{-2\bar{k}_ns} ds \geq 2\sigma_n^2 \kappa (1 - e^{-t/\kappa})$ and, since $|\delta_n| \geq \text{Re } \delta_n = \kappa^{-1} - k_n \geq \kappa^{-1}$, $\mathbb{E}^{\text{mid}} Y_n \geq 2\kappa^{-1} [T - \kappa(1 - e^{-T/\kappa})] > 0$ on infinitely many blocks; the necessity argument of (b) then gives $\sum_n (1 - \mathfrak{h}_n) = \infty$ and singularity on Ω . The product laws exist on Ω whatever the signs of the k_n , but their coordinate variances lack uniform dissipative control: they grow linearly in t when $k_n = 0$ and exponentially when $k_n < 0$. Consequently $\sum_n \omega_n^2 |x_n|^2$ need not

converge.

(d) Substitute $p(z_n) = -\lambda_n$ and sum $\lambda_n = (1 + \kappa^2 \omega_n^2)/(2\kappa)$ over $n \leq N$; the sum diverges, so (b) gives singularity in the limit. For the affinity write (22) as

$$\mathfrak{h}_n = \frac{2e^{(\bar{k}_n - \mathfrak{m}_n)T}}{(1 + \bar{k}_n/\mathfrak{m}_n) + (1 - \bar{k}_n/\mathfrak{m}_n)e^{-2\mathfrak{m}_n T}},$$

whose denominator lies between 1 and 2 because $0 < \bar{k}_n \leq \mathfrak{m}_n$; hence $e^{(\bar{k}_n - \mathfrak{m}_n)T} \leq \mathfrak{h}_n \leq 2e^{(\bar{k}_n - \mathfrak{m}_n)T}$. Now $\mathfrak{m}_n - \bar{k}_n = |\delta_n|^2/(4(\mathfrak{m}_n + \bar{k}_n))$, with $|\delta_n|^2 = \lambda_n^2$ and $\bar{k}_n = \frac{1}{2}(\lambda_n + \kappa^{-1})$ for the model. Since $\mathfrak{m}_n + \bar{k}_n \geq 2\bar{k}_n \geq \lambda_n$, one has $\mathfrak{m}_n - \bar{k}_n \leq \lambda_n/4$, whence $\prod_{n \leq N} \mathfrak{h}_n \geq e^{-T\mathcal{E}_N/4}$. In the other direction, $\mathfrak{m}_n^2 = \frac{1}{2}\lambda_n^2 + \frac{\lambda_n}{2\kappa} + \frac{1}{4\kappa^2} \geq (\frac{\lambda_n}{\sqrt{2}} + \frac{1}{2\sqrt{2}\kappa})^2$ gives $\bar{k}_n - \mathfrak{m}_n \leq -c_1\lambda_n + (1 - \frac{1}{\sqrt{2}})\frac{1}{2\kappa}$ and the displayed upper bound. Finally $(\mathfrak{m}_n + \bar{k}_n)/\lambda_n \rightarrow \frac{1}{\sqrt{2}} + \frac{1}{2}$, so $(\mathfrak{m}_n - \bar{k}_n)/\lambda_n \rightarrow (4(\frac{1}{\sqrt{2}} + \frac{1}{2}))^{-1} = c_1$; as $-\log \mathfrak{h}_n$ differs from $(\mathfrak{m}_n - \bar{k}_n)T$ by at most $\log 2$ and $\mathcal{E}_N \asymp N^3$, it follows that $-\log \prod_{n \leq N} \mathfrak{h}_n \sim c_1 T \mathcal{E}_N$. $\square$

Remark 4.12 (The price, its invariant, and the sequence of tilted markets). (i) At the reference calibration $(\kappa, \theta^*) = (4, 30)$ and $T = 1$ the truncated model has

N	5	10	16	32
$\mathcal{E}_N$	1.83	9.69	34.8	255
$\sqrt{\mathcal{E}_N}$	1.35	3.11	5.90	16.0
$\prod_{n \leq N} \mathfrak{h}_n$	0.92	0.37	$6 \cdot 10^{-3}$	$2 \cdot 10^{-21}$

(ii) At each finite order, $|\gamma_t|$ bounds the instantaneous Sharpe ratio of every self-financing portfolio in the N -block market: a portfolio with volatility vector v has excess drift $v \cdot \gamma$, hence Sharpe ratio $v \cdot \gamma/|v| \leq |\gamma|$, with equality when $v \parallel \gamma$. A panel of $2N$ bonds as in (iv) attains the equality at all but finitely many times of its horizon, since its volatility vectors then span $\mathbb{R}^{2N}$. Thus $\sqrt{\mathcal{E}_N}$ is the stationary root-mean-square of the market price of risk $|\gamma_t|$, equivalently of the maximal instantaneous Sharpe ratio of the N th market — a continuous-time local analogue of the Hansen–Jagannathan bound [23] — and $\sqrt{\mathcal{E}} = \lim_N \sqrt{\mathcal{E}_N}$ is its limit. Unlike $\mathfrak{d}_N$ of Proposition 4.8 (Remark 4.9), $\mathcal{E}_N$ is a market-price-of-risk energy, precisely because under the isotropic loading the defect lies in the range of the volatility. (iii) The invariant is coordinate-free. For a factor generator K with noise loading Σ , possibly rectangular, whose covariance $C = \Sigma \Sigma^\top$ is invertible, and with stationary covariance V , the minimum-norm tilt carrying the induced dynamics to the class is $\gamma_{\min} = \Sigma^\top C^{-1} \Xi x$ with $\Xi = K - M^\top$, and its stationary energy is $\mathbb{E}|\gamma_{\min}|^2 = \text{tr}(C^{-1} \Xi V \Xi^\top) = \|C^{-1/2} \Xi V^{1/2}\|_{\text{HS}}^2$. Finiteness of this quantity is a property of the operator and the noise, invariant under changes of factor coordinates; in the block-diagonal isotropic setting it reduces to $\mathcal{E} < \infty$, with a value independent of the intensities. (iv) Fix a trading horizon $H \in (0, \theta^*)$ and take $T = H$ in (a)–(d). At each order take the market of the bank account and $2N$ bonds with fixed maturities $T_j \in (H, \theta^*)$, so that every bond is alive on $[0, H]$ with time to maturity in the modelled range, and let the deterministic sector be transported as in Proposition 2.7. The $2N \times 2N$ volatility matrix has entries $\sigma_n I_n^s(T_j - t)$ and $\sigma_n I_n^c(T_j - t)$, $t \in [0, H]$. The $2N$ functions I_n^s, I_n^c ($n \leq N$) are linearly independent real-analytic functions on every open interval, since their derivatives are the distinct characters b_n^s, b_n^c ; so by the evaluation argument in the proof of Theorem 4.6 the matrix is invertible at $t = 0$ for an open dense set of maturity tuples in $(H, \theta^*)^{2N}$, and, being real-analytic in t , it is invertible for all but

finitely many $t \in [0, H]$, hence dt -a.e. On the natural filtration of the $2N$ Brownian motions the market is therefore complete, with Q_N as its unique pricing measure. In the singular case $\prod_{n \leq N} \mathfrak{h}_n \rightarrow 0$ by Kakutani's criterion, so the sequences (P_N) and (Q_N) are entirely separated; that separation is the criterion for strong asymptotic arbitrage in the sense of [29] (see also [38]). For the model this holds for every $H \in (0, \theta^*)$ by (d), and it is a rigorous finite-sequence form of the infinite statement left open below.

The diagnosis is now complete at every finite order. *Every nontrivial finite sine truncation of the envelope, equipped with the induced drift, admits no equivalent local martingale measure;* once the silent coordinates are loaded, every truncation admits one, at a price that diverges with the truncation order. In the infinite model the factor-law statement is settled on the loaded side, where the laws are mutually singular (Proposition 4.11(b)). The statement for the traded bonds is the open problem of Remark 3.4: it presupposes traded-bond processes that are semimartingales under the cylindrical dynamics. The next section stops asking whether the model is arbitrage-free and constructs the arbitrage-free model that its envelope and its noise admit.

5 The arbitrage-free completion and its singular limit

Section 4 returned a verdict. At every finite sine truncation the deformation drift lies outside its volatility span, the defect is realisable in traded bonds, and it does not fade under refinement; once the silent coordinates are loaded, it carries a price that diverges. This section does what the conclusion of [13] asks and turns the verdict into a construction — keep the model's basis envelope and its noise, repair the drift, and take $N \rightarrow \infty$ to see where the singularity of maturity-white risk reappears. The repair is exact at every N . What it costs, and at which level its limit lives, are the two results here: a short-end drift discontinuity with an explicit boundary layer, and a limit for the convexity and the bond-return volatilities but none for the curve. Section 6 then asks whether that cost survives the model's own stochastic spread.

5.1 The minimal completion and the exact limit

The repair is a definition: keep the envelope of Proposition 4.1 and the sine-only noise, and impose the drift that Proposition 2.7 requires.

Definition 5.1 (Completion). Fix $\kappa, \theta^*, \sigma_0 > 0$ and $N \geq 1$. The N -block completion is the model of Assumption 2.6 with basis $\mathbf{B} = \bigoplus_{n \leq N} \sqrt{\frac{2}{\theta^*}} e^{-\tau/\kappa} (\sin \omega_n \tau, \cos \omega_n \tau)^\top$, generating matrix $M = \bigoplus_n M_n$, factor drift $K = M^\top$, noise loading only on the sine coordinates with intensity σ_0 (so $C = \bigoplus_n \sigma_0^2 \text{diag}(1, 0)$ — exactly the model's sources of randomness), and φ_t transported by (6) with $q_t^N = q^N$ below. By Proposition 2.7 each completion is arbitrage-free.

The cosine coordinates carry no noise but are excited through the rotation in $M^\top$; the completion is a $2N$ -dimensional Gaussian factor HJM model. Its convexity is $q^N(\tau) = \mathbf{B}(\tau)^\top C \mathbf{I}(\tau) = \sum_{n \leq N} q_n(\tau)$. Per block only the sine coordinate is loaded, so $q_n(\tau) = \sigma_0^2 b_n^s(\tau) \int_0^\tau b_n^s(u) du$ with $b_n^s(\tau) = \sqrt{2/\theta^*} e^{-\tau/\kappa} \sin \omega_n \tau$, and the standard exponential-trigonometric integral gives the closed form

$$q_n(\tau) = \frac{2\sigma_0^2}{\theta^*} \cdot \frac{e^{-\tau/\kappa}}{\kappa^{-2} + \omega_n^2} \left[\omega_n \sin \omega_n \tau - \frac{e^{-\tau/\kappa}}{\kappa} \sin^2 \omega_n \tau - \omega_n e^{-\tau/\kappa} \sin \omega_n \tau \cos \omega_n \tau \right]. \quad (23)$$

Theorem 5.2 (Exact completion limit). *For every $\kappa, \theta^*, \sigma_0 > 0$ and every $\tau \in (0, \theta^*)$,*

$$\boxed{\lim_{N \rightarrow \infty} q^N(\tau) = q^\infty(\tau) = \frac{\sigma_0^2}{2} e^{-2\tau/\kappa}} \quad (24)$$

— independent of θ^* — with convergence uniform on compact subsets of $(0, \theta^*)$. Meanwhile $q^N(0) = 0$ for every N (as any HJM drift must satisfy at $\tau = 0$), so the limit has a jump $q^\infty(0^+) = \sigma_0^2/2$: the convergence develops a Gibbs boundary layer at the short end with universal profile

$$q^N\left(\frac{u\theta^*}{N\pi}\right) \longrightarrow \sigma_0^2 G(u), \quad G(u) := \frac{2}{\pi} \left[\text{Si}(u) - \frac{1}{2} \text{Si}(2u) \right], \quad (25)$$

uniformly on compacts of $[0, \infty)$. The limiting profile has a strict first local maximum at $u = \pi$, of height

$$\sigma_0^2 G(\pi) = \frac{2}{\pi} \left[\text{Si}(\pi) - \frac{1}{2} \text{Si}(2\pi) \right] \sigma_0^2 = 0.727568 \sigma_0^2 = 1.45514 \times q^\infty(0^+),$$

in the following precise local sense: for fixed $0 < a < \pi < b < 2\pi$, every maximiser u_N of $u \mapsto q^N(u\theta^*/(N\pi))$ on $[a, b]$ satisfies $u_N \rightarrow \pi$, and its value converges to $\sigma_0^2 G(\pi)$.

Proof. Write $\Delta_n = \kappa^{-2} + \omega_n^2$, use $\sin^2 = \frac{1}{2}(1 - \cos 2\omega_n\tau)$, $\sin \cos = \frac{1}{2} \sin 2\omega_n\tau$, and sum (23) against the three classical series (Fourier expansions of the sawtooth and of hyperbolic functions on the lattice $\omega_n = n\pi/\theta^*$; e.g. Gradshteyn–Ryzhik 1.441–1.445):

$$S_1(x) := \sum_{n \geq 1} \frac{\omega_n \sin \omega_n x}{\Delta_n} = \frac{\theta^* \sinh \frac{\theta^* - x}{\kappa}}{2 \sinh \frac{\theta^*}{\kappa}}, \quad S_2(x) := \sum_{n \geq 1} \frac{\cos \omega_n x}{\Delta_n} = \frac{\kappa \theta^* \cosh \frac{\theta^* - x}{\kappa}}{2 \sinh \frac{\theta^*}{\kappa}} - \frac{\kappa^2}{2},$$

valid for $0 < x < 2\theta^*$ (resp. $0 \leq x \leq 2\theta^*$), and $S_3 := S_2(0)$. Then

$$q^\infty(\tau) = \frac{2\sigma_0^2}{\theta^*} e^{-\tau/\kappa} \left[S_1(\tau) - \frac{e^{-\tau/\kappa}}{2\kappa} (S_3 - S_2(2\tau)) - \frac{e^{-\tau/\kappa}}{2} S_1(2\tau) \right],$$

the doubled arguments requiring $\tau < \theta^*$. Elementary hyperbolic algebra (set $u = \tau/\kappa$, $y = \theta^*/\kappa$; the sum telescopes to $[e^{y-2u} - e^{-y-2u}]/(4 \sinh y)$, whose numerator factors as $2e^{-2u} \sinh y$, leaving $\frac{1}{2}e^{-2u}$) collapses every θ^* -dependent term and yields (24). Convergence: after the trigonometric reduction the tails are of two kinds. The absolutely convergent ones have $O(\Delta_n^{-1}) = O(n^{-2})$ coefficients. The conditionally convergent ones have coefficients ω_n/Δ_n , monotone $\rightarrow 0$ after finitely many terms, against $\sin \omega_n x$, whose partial sums are bounded uniformly for x in compact subsets of $(0, 2\theta^*)$; Dirichlet's test then gives uniform convergence on compacts of $(0, \theta^*)$. At $\tau = 0$ every term of (23) vanishes; at $\tau = \theta^*$ every $\sin \omega_n \theta^* = 0$, so $q^N(\theta^*) = 0 \neq q^\infty(\theta^{*-})$.

(Layer profile.) Put

$$a := \kappa^{-1}, \quad h := \frac{\pi}{\theta^*}, \quad x := h\tau = \frac{u}{N};$$

thus $\omega_n = hn$ and $\tau = u/(Nh)$. Fix $U < \infty$ and let $0 \leq u \leq U$. The elementary bounds

$$|\sin(nx)| \leq \frac{nU}{N}, \quad |\sin(2nx)| \leq \frac{2nU}{N},$$

give

$$\begin{aligned} \sum_{n \leq N} \left| \frac{\omega_n}{a^2 + \omega_n^2} - \frac{1}{\omega_n} \right| (|\sin(nx)| + |\sin(2nx)|) \\ \leq \frac{3a^2U}{h^3N} \sum_{n \geq 1} \frac{1}{n^2} = O_U(N^{-1}), \\ \sum_{n \leq N} \frac{\sin^2(nx)}{a^2 + \omega_n^2} \leq \frac{U^2}{h^2N}. \end{aligned}$$

Here we used

$$\left| \frac{\omega_n}{a^2 + \omega_n^2} - \frac{1}{\omega_n} \right| = \frac{a^2}{\omega_n(a^2 + \omega_n^2)} \leq \frac{a^2}{h^3n^3}.$$

Moreover,

$$|e^{-a\tau} - 1| \leq a\tau \leq \frac{aU}{hN}, \quad \sum_{n \leq N} \frac{|\sin(nx)| + |\sin(2nx)|}{n} \leq 3U.$$

Applying these estimates term by term to (23) yields, uniformly for $u \in [0, U]$,

$$q^N\left(\frac{u\theta^*}{N\pi}\right) = \sigma_0^2 \frac{2}{\pi} \sum_{n \leq N} \frac{\sin(nu/N) - \frac{1}{2}\sin(2nu/N)}{n} + O_U(N^{-1}). \quad (26)$$

For

$$g_u(y) := \begin{cases} \sin(uy)/y, & y > 0, \\ u, & y = 0, \end{cases}$$

the family $\{g_u : 0 \leq u \leq U\}$ has uniformly bounded first derivatives on $[0, 1]$. Hence its right Riemann sums satisfy

$$\sum_{n \leq N} \frac{\sin(nu/N)}{n} = \frac{1}{N} \sum_{n \leq N} g_u(n/N) = \int_0^1 g_u(y) dy + O_U(N^{-1}) = \text{Si}(u) + O_U(N^{-1}),$$

uniformly for $u \in [0, U]$; applying the same argument with $2u$ gives

$$\sum_{n \leq N} \frac{\sin(2nu/N)}{n} = \text{Si}(2u) + O_U(N^{-1}).$$

Substitution in (26) proves

$$q^N\left(\frac{u\theta^*}{N\pi}\right) \longrightarrow \sigma_0^2 \frac{2}{\pi} \left[\text{Si}(u) - \frac{1}{2} \text{Si}(2u) \right]$$

uniformly on compact subsets of $[0, \infty)$. Since $G'(u) = \frac{2}{\pi} \frac{\sin u (1 - \cos u)}{u}$ is positive on $(0, \pi)$ and negative on $(\pi, 2\pi)$, $u = \pi$ is the unique maximiser on every fixed interval $[a, b]$ as in the statement. Uniform convergence and the standard argmax lemma give the asserted convergence of the window-restricted maximisers and values; the conclusion is local to the window $[a, b]$ of the statement. $\square$

The θ^* -independence of the short-end value has a one-line explanation. In the layer the leading sums are the sawtooth series $\sum_{n \geq 1} \sin(nx)/n = \frac{\pi-x}{2}$ and $\sum_{n \geq 1} \sin(2nx)/n = \frac{\pi-2x}{2}$ for $x \in (0, \pi)$, and $\frac{2}{\pi} \left[\frac{\pi-x}{2} - \frac{\pi-2x}{4} \right] = \frac{1}{2}$ *identically*: the x -dependence cancels, which is the sawtooth form of the θ^* -independence of $q^\infty(0^+)$.

The two constants deserve comment. The relative first-layer peak 1.45514 exceeds the classical Wilbraham–Gibbs value $\frac{2}{\pi} \text{Si}(\pi) = 1.17898$ for a plain jump. The reason is that $\text{Si}(2\pi)$ is the first local *minimum* of Si: the second harmonic reaches its maximal downward excursion exactly where the first peaks, and the two excursions reinforce rather than cancel. And a corresponding far-end discontinuity is recorded in Remark 6.6: the two one-sided limits at $\tau = \theta^*$ are $\pm \frac{\sigma_0^2}{2} e^{-2\theta^*/\kappa}$ — a full jump of size $\sigma_0^2 e^{-2\theta^*/\kappa}$ — while $q^N(\theta^*) = 0$ for every N .

Remark 5.3 (Interpretation: the drift cost of the minimal completion). Three readings of Theorem 5.2. (i) *Kennedy’s half-delta is exact at finite span*. For a random-field HJM with spatial covariance kernel $c(\tau_1, \tau_2)$, the formal drift condition is $\alpha(\tau) = \int_0^\tau c(\tau, u) du$ [33, 22, 39]. The cylindrical noise has $c(\tau, u) = \sigma_0^2 e^{-2u/\kappa} \delta(\tau - u)$ distributionally (completeness of (8) in the weighted space), and integrating a boundary delta with half its mass gives exactly $\frac{\sigma_0^2}{2} e^{-2\tau/\kappa}$. The completion theorem turns this formal computation into a rigorous limit statement — and shows the finite- θ^* corrections cancel *identically*, which the heuristic does not predict. (ii) *The jump is the shadow of Proposition 3.3*. The same white-in-maturity noise that destroys pointwise semimartingality shows up, after finite-dimensional repair and passage to the limit, as an irremovable $\sigma_0^2/2$ discontinuity of the short-end drift. Arbitrage-free approximants exist at every N , but their drifts develop a boundary layer of width θ^*/N rather than converging smoothly. Within this selected completion, maturity-local volatility therefore entails a singular short-end HJM drift cost. (iii) *Non-uniqueness of the completion is quantifiable*. Loading noise isotropically on both coordinates ($C = \sigma_0^2 \mathbf{1}_2$ per block) gives instead, again as the first-span interior limit on $0 < \tau < \theta^*$,

$$q_\infty^{\text{iso}}(\tau) = \sigma_0^2 \left[e^{-2\tau/\kappa} + \frac{\kappa}{\theta^*} \left(e^{-2\tau/\kappa} - e^{-\tau/\kappa} \right) \right].$$

Indeed, adding the cosine-loaded analogue of (23) to the sine-loaded term and using the same series S_1, S_2, S_3 as in the proof of Theorem 5.2 gives

$$q_\infty^{\text{iso}}(\tau) = \frac{2\sigma_0^2}{\theta^*} \left\{ e^{-\tau/\kappa} [S_1(\tau) + \kappa^{-1} S_2(\tau)] - \kappa^{-1} e^{-2\tau/\kappa} S_3 \right\}.$$

The identities

$$S_1(\tau) + \kappa^{-1} S_2(\tau) = \frac{\theta^* e^{(\theta^* - \tau)/\kappa}}{2 \sinh(\theta^*/\kappa)} - \frac{\kappa}{2}, \quad \kappa^{-1} S_3 = \frac{\theta^*}{2} \coth(\theta^*/\kappa) - \frac{\kappa}{2}$$

reduce this expression to the displayed formula. It has the doubled short-end jump σ_0^2 , a genuine $O(\kappa/\theta^*)$ span correction, and changes sign when $\tau > \kappa \log(1 + \theta^*/\kappa)$. The completion is a choice; the sine-only one is canonical in that it adds no sources of randomness to [13]. The layer profile is likewise a property of the sharp sine cutoff of Definition 5.1: smoothed truncations (Fejér and its relatives) damp or remove the overshoot, so the Gibbs constant belongs to the completion choice, not to the limit.

5.2 No curve-valued limit

The convexity density converges; the curve laws do not, and not merely in stationarity. The next statement makes this quantitative, and identifies what does survive: the bond-level functionals.

Proposition 5.4 (No curve-level limit for the equal-intensity completion). *Start each completion from a deterministic factor state, or more generally from an initial state independent of its future Brownian noise. For every fixed $t > 0$ and block n , the covariance contributed by the stochastic convolution is*

$$V_n(t) = \int_0^t e^{K_n s} C_n e^{K_n^\top s} ds = \frac{\sigma_0^2 \kappa}{4} (1 - e^{-2t/\kappa}) \mathbf{1}_2 + R_n(t), \quad \|R_n(t)\| = O_t(\omega_n^{-1}).$$

Consequently, at every fixed $\tau \geq 0$, the stochastic-convolution contribution satisfies

$$\text{Var}_{\text{conv}} f_t^N(\tau) = \frac{\sigma_0^2 \kappa}{2\theta^*} (1 - e^{-2t/\kappa}) N e^{-2\tau/\kappa} + O_t(\log N) \longrightarrow \infty,$$

and, whenever the initial state has a finite second moment, the total variance is at least this quantity, with equality for deterministic initial states. The non-tightness conclusion itself needs no moment assumption — only independence. The point-evaluation laws are therefore not tight, and neither are the curve laws in H_w , since evaluation is continuous there. In stationarity the same formulas hold with $(1 - e^{-2t/\kappa})$ replaced by 1. Bond-level functionals survive: $\|\Sigma(\tau)\|^2 = |C^{1/2} \mathbf{B}_I(\tau)|^2$ is N -uniformly finite.

Proof. Since $K_n = -\kappa^{-1} \mathbf{1}_2$ plus a skew-symmetric rotation generator, $e^{K_n s} C_n e^{K_n^\top s}$ is $e^{-2s/\kappa}$ times a rotation of $C_n = \sigma_0^2 \text{diag}(1, 0)$. Its nonoscillatory part is $\frac{\sigma_0^2}{2} e^{-2s/\kappa} \mathbf{1}_2$, whose integral gives the displayed scalar term; integration of the sine and cosine terms of frequency $2\omega_n$ gives $\|R_n(t)\| = O_t(\omega_n^{-1})$. Now $\mathbf{B}_n(\tau)^\top \mathbf{1}_2 \mathbf{B}_n(\tau) = (2/\theta^*) e^{-2\tau/\kappa}$ and $\sum_{n \leq N} \omega_n^{-1} = O(\log N)$. Independence of the initial state and future noise makes covariances additive, proving the variance bound. For the bond-level claim, $I_n^s(\tau) = O(\omega_n^{-1})$ uniformly on compacts, so $\sum_n |I_n^s(\tau)|^2 < \infty$ uniformly in N . Finally, convolution with a centred Gaussian of variance tending to infinity has concentration on every interval of fixed length tending to zero; therefore no choice of independent initial law can restore pointwise tightness. $\square$

The mechanism is structural. The arbitrage-free repair replaces the model's decay $\lambda_n \asymp n^2$ by the uniform $1/\kappa$, and with it goes the mode-variance summability that made the cylindrical field function-valued (Proposition 3.2). What converges is the convexity density (Theorem 5.2) and the bond-return volatility structure: the limit lives at the level of bonds, not of the curve. A distribution-valued or bond-level limit model — traded-bond semimartingales, an equivalent martingale measure, cross-maturity consistency — is the natural next object, and is the first open problem of §7.3. Thus, *within this equal-amplitude cylindrical sine refinement and its minimal finite-dimensional HJM completion*, the calculation exhibits a trilemma: retaining the maturity-local noise, imposing no-arbitrage at every truncation, and obtaining a function-valued refinement limit — any two. For the completion, factor count is therefore economically consequential rather than a numerical grid parameter: bond-return volatility functionals remain controlled, whereas instantaneous forward-curve laws do not stabilize under refinement.

The third corner of the trilemma is not empty, and Proposition 4.11 says what lives there. Give up the maturity-local noise for the isotropic trace-class loading of §4.7. Then the factor




Figure 2: (a) Factor autocorrelation on block $n = 2$ ($\kappa = 4$, $\theta^* = 30$): the arbitrage-free dynamics $e^{M_n^\top h}$ spirals (decay $1/\kappa$, rotation ω_n), the model's $e^{-\lambda_n h}$ decays along a ray. This is Corollary 4.5 in a picture. (b) The completion's convexity: partial sums q^N against the exact limit $\frac{\sigma_0^2}{2} e^{-2\tau/\kappa}$; every q^N vanishes at $\tau = 0$ while the limit jumps to $\sigma_0^2/2$; the Gibbs layer narrows as θ^*/N . The layer has the universal profile $\sigma_0^2 G(N\pi\tau/\theta^*)$, $G(u) = \frac{2}{\pi}[\text{Si}(u) - \frac{1}{2}\text{Si}(2u)]$, with reinforced first-layer peak $1.455 \times \sigma_0^2/2$ near $\tau = \theta^*/N$ (Theorem 5.2). (c) The two arbitrage-free completions of Remark 5.3(iii).

drift $M^\top$ of the completion carries an H_w -valued curve, no-arbitrage holds at every resolution, and the question becomes how much of the model's curvature term can be kept in the historical drift at a finite price. Let $\mathfrak{D} := \partial_\tau + \kappa^{-1}$, which acts on the block V_n as $\omega_n J$, and for a scale $\nu > 0$ (in years) consider the *mollified* generator, as a formal expression on quasi-exponentials,

$$\tilde{A}_\nu := \partial_\tau + \frac{\kappa}{2} \partial_\tau^2 e^{\frac{\nu^2}{2} \mathfrak{D}^2}, \quad p_\nu(z_n) = z_n + \varpi_n(p(z_n) - z_n), \quad \varpi_n := e^{-\nu^2 \omega_n^2/2},$$

with entire symbol $p_\nu(z) = z + \frac{\kappa}{2} z^2 e^{\frac{\nu^2}{2} (z + \kappa^{-1})^2}$. The definition that carries the mathematics is blockwise: $K_\nu := \bigoplus_n p_\nu(M_n)^\top$, and by the block calculus of §4.7 $p_\nu(M_n)^\top = (1 - \varpi_n) M_n^\top + \varpi_n p(M_n)^\top$, the convex combination of transport and the model's drift. The multiplier has a reading in the model's own terms: since $\frac{\kappa}{2} \mathfrak{D}^2 = \tilde{A} + \frac{1}{2\kappa}$ as formal expressions, $e^{\frac{\nu^2}{2} \mathfrak{D}^2}$ is $e^{\nu^2/(2\kappa^2)}$ times the formal flow $e^{(\nu^2/\kappa)\tilde{A}}$, so the curvature penalty is applied to the curve after the cylindrical flow has acted on it for ν^2/κ years. This is the formal calculus of §3 on the doubled envelope, not the closed Dirichlet semigroup, which does not preserve the cosine partners; on the sine eigenspaces the two agree. In the notation of (21), $|\delta_n(\nu)| = \varpi_n \lambda_n$ and $k_n(\nu) = \kappa^{-1} - d_n \varpi_n$, with $k_n(\nu) - \frac{1}{2\kappa} = \frac{1 - \varpi_n}{2\kappa} + \frac{\nu^2}{2} \omega_n^2 \varpi_n > 0$, so

$$\mathcal{E}_N(\nu) := \sum_{n \leq N} \frac{\varpi_n^2 \lambda_n^2}{\kappa^{-1} - d_n \varpi_n}, \quad \mathcal{E}(\nu) := \lim_{N \rightarrow \infty} \mathcal{E}_N(\nu) < \infty \quad (\nu > 0),$$

and

$$\mathcal{E}(\nu) \sim \frac{\kappa \theta^*}{2\sqrt{2\pi}} \nu^{-3} \quad (\nu \rightarrow 0).$$

For the asymptotic, write $k_n(\nu) = a_n + u_n$ with $a_n := (1 - \varpi_n)/\kappa$ and $u_n := \varpi_n \lambda_n$; then $\mathfrak{e}_n = u_n^2/(a_n + u_n)$ and $0 \leq u_n - \mathfrak{e}_n \leq \min(a_n, u_n)$. The main term $\sum_n u_n = \frac{\kappa}{2} \sum_n \omega_n^2 \varpi_n + \frac{1}{2\kappa} \sum_n \varpi_n$ is, up to $O(\nu^{-2})$, the Riemann sum with spacing π/θ^* of the unimodal function $\frac{\kappa}{2} \omega^2 e^{-\nu^2 \omega^2/2}$,

whose integral is $\frac{\kappa}{2}\sqrt{\pi/2}\nu^{-3}$, plus an $O(\nu^{-1})$ term. The remainder is $o(\nu^{-3})$: splitting at $\omega = L/\nu$, the blocks below contribute at most $\sum a_n \leq \sum \nu^2 \omega_n^2 / (2\kappa) = O(L^3 \nu^{-1})$, and those above at most $\sum u_n$, which is ν^{-3} times a Gaussian tail integral that vanishes as $L \rightarrow \infty$. By Proposition 4.11, for every $\nu > 0$ the infinite factor laws are equivalent. The mollified curvature term is carried entirely by the historical drift, as an essentially affine risk premium in the sense of Remark 2.10; under the pricing measure the factor drift is the completion's $M^\top$. The symbol test of Theorem 4.3 still fails on every block: the defect is the model's, scaled by ϖ_n . Its price is $\sqrt{\mathcal{E}_N(\nu)}$, the stationary root-mean-square of the maximal Sharpe ratio the premium commands in the N th truncated market, with $\sqrt{\mathcal{E}(\nu)}$ its limit (Remark 4.12(ii)). What is bought is modest and exactly measurable: the fraction ϖ_n of the curvature term kept on block n . At the reference calibration a Sharpe ratio of 0.4 corresponds to $\nu \approx 5.7$ years, where ϖ_n is 0.84, 0.49 and 0.20 on the 60-, 30- and 20-year modes, 0.06 on the 15-year mode and below 0.02 beyond: the curvature penalty survives on the longest wavelengths only. The Gaussian damping is one choice among many: within the block-diagonal class of Proposition 4.11, any blockwise deviation from transport with $\sum_n |\delta_n|^2 / k_n < \infty$ is admissible, and only the ν^{-3} asymptotic and the calibration above are specific to $\varpi_n = e^{-\nu^2 \omega_n^2 / 2}$. Keeping the model's drift exactly on the two longest modes and transport elsewhere costs as much as $\nu = 4$ (both give a Sharpe ratio of 0.60). What this corner gives up is the cylindrical noise itself: with $\sigma_n \equiv \sigma_0$ the stationary variances $\sigma_0^2 / (2k_n)$ do not sum. The short-end cost of Theorem 5.2 therefore does not arise here, and the case for one shock per maturity that motivates the model is met no better than by the completion (§7.3).

5.3 The long end

One question remains before the spread is restored. In any arbitrage-free model the long zero-coupon rate, where it exists as a limit, is almost surely nondecreasing [16]; the general rigorous proof is [27]. Does complexification interact with this long-maturity rigidity of Gaussian HJM models? It does not, in either direction.

Proposition 5.5 (DIR and complexification are orthogonal). *In a model of Assumption 2.6:*

- (i) *every stable oscillatory block satisfies $|\mathbf{B}_b(\tau)| \leq Ce^{-\delta\tau} \rightarrow 0$ for some $\delta > 0$ and has bounded $\mathbf{I}_b(\tau)$, so $\frac{1}{\tau}\mathbf{I}_b(\tau)^\top \mathbf{X}_{b,t} \rightarrow 0$: complex modes contribute nothing to the long yield;*
- (ii) *with no level mode ($\ker M = 0$) the whole factor contribution to the long yield vanishes;*
- (iii) *a level mode with $C_t e_0 \neq 0$ — volatility-loaded or cross-loaded, the same condition — has $\mathbf{I}_0(\tau) = b_0 \tau$ and forces q_t to grow linearly with slope $b_0^\top e_0^\top C_t e_0$, the standard Gaussian-HJM long-end pathology; a level mode with $C_t e_0 \equiv 0$ is frozen, and so is its contribution to the long yield.*

The long-end behaviour thus depends only on the (generalised) kernel of M and on how volatility loads it: complexification neither causes nor cures anything at the long end.

Proof. For the model class, consistency holds for every initial state, so $\mathbf{B}^\top x = r_0 - \varphi_0 \in H_w$ for every x , i.e. every component of $\mathbf{B}$ lies in H_w . Controllability makes $\mathbf{B}(0)$ cyclic for M (Lemma 2.4), so each Jordan cell of M contributes its top chain function to the components of $\mathbf{B}$; H_w -membership (§2.1) then excludes exponents with $\operatorname{Re} \lambda > 0$, undamped oscillations, and any Jordan tail over the kernel. Hence, in real Jordan coordinates, $\mathbf{B}(\tau) = b_0 e_0 + \mathbf{B}_{\text{dec}}(\tau)$:

an at most one-dimensional constant level loading on $\ker M = \ker_{\text{gen}} M$, and a strictly stable remainder with $|\mathbf{B}_{\text{dec}}(\tau)| \leq Ce^{-\delta\tau}$ for some $\delta > 0$.

(i) *Oscillatory blocks.* They lie in $\mathbf{B}_{\text{dec}}$, and M_b is invertible ($\omega \neq 0$), so $\mathbf{I}_b(\tau) = M_b^{-1}(e^{M_b\tau} - \mathbf{1})\mathbf{B}_b(0)$ is bounded; for every fixed t , $\frac{1}{\tau}\mathbf{I}_b(\tau)^\top \mathbf{X}_{b,t} \rightarrow 0$ surely. Complex modes contribute nothing to any long-maturity average, in particular not to the long yield.

(ii) *No level mode.* If $\ker M = 0$, the bound applies to all of $\mathbf{B}$: $\mathbf{I}$ is bounded and the factor contribution $\frac{1}{\tau}\mathbf{I}(\tau)^\top \mathbf{X}_t$ vanishes in the limit.

(iii) *Level mode.* With $\mathbf{I}(\tau) = \tau b_0 e_0 + \mathbf{I}_{\text{dec}}(\tau)$, $\mathbf{I}_{\text{dec}} := \int_0^\tau \mathbf{B}_{\text{dec}}$ bounded, and $\mathbf{B}(\tau) = b_0 e_0 + O(e^{-\delta\tau})$,

$$q_t(\tau) = \mathbf{B}(\tau)^\top C_t \mathbf{I}(\tau) = \tau b_0^2 e_0^\top C_t e_0 + o(\tau).$$

Since $C_t \succeq 0$, $e_0^\top C_t e_0 = |\Sigma_t^\top e_0|^2$ vanishes exactly when $C_t e_0 = 0$: volatility-loaded and cross-loaded are the same condition, and there is no third case. If $\int_0^t e_0^\top C_s e_0 ds > 0$, (6) gives φ_t a term linear in τ with that integral times b_0^2 as slope. The curve then leaves H_w (nonconstant affine tails have infinite norm, Theorem 4.3(b)), and the long yield $\lim_{\tau} \frac{1}{\tau} \int_0^\tau f_t$ is $+\infty$ for such t : the standard Gaussian-HJM long-end explosion, of which the Ho–Lee check (Remark 2.9) is the $n = 1$ instance. If instead $C_t e_0 \equiv 0$, then also $e_0^\top C_t \equiv 0$, every term of q_t containing the level direction dies, and $q_t(\tau) = \mathbf{B}_{\text{dec}}(\tau)^\top C_t \mathbf{I}_{\text{dec}}(\tau) = O(e^{-\delta\tau})$. Moreover the level factor is constant in time ($dx_0 = e_0^\top M^\top \mathbf{X} dt + e_0^\top \Sigma_t d\mathbf{W} = 0$ because $Me_0 = 0$ and $\Sigma_t^\top e_0 = 0$), so the long yield exists and is frozen: trivially consistent with [16], with nothing left for the pathology to act on.

Every statement above depends only on $\ker M$ and on $C_t e_0$; no oscillatory quantity enters any of them. That is the orthogonality claim. $\square$

Remark 5.6. The finite θ_{max} of [13] is a deliberate choice made precisely to sidestep the $\theta \rightarrow \infty$ regime, with [16] cited; on a compact maturity interval the theorem does not apply. The accurate statement is conditional: the stationary Gaussian structure does not extend to $\theta_{\text{max}} = \infty$, and [13] calls the limit “not innocent.”

That completes the analysis at fixed spread intensity: the model is repaired at every N , the cost of the repair — its HJM convexity, in the sense of Remark 5.3 — is computed exactly, and the refinement limit is characterised and shown not to be a curve. One feature has been held fixed throughout — the spread that multiplies the deformation — and the next section restores it, to see whether the cost survives.

6 The cost survives the stochastic spread

The model multiplies its deformation by a stochastic spread s_t ; §§4–5 held that spread fixed, and the obvious objection to Theorem 5.2 is that its exact cost is an artefact of freezing it. This section restores the spread and shows that it is not. The algebra closes in two conventions. In the *affine* convention the squared intensity $v_t = s_t^2$ is a CIR factor and the deformation is conditionally Gaussian (§§6.1–6.4). There the short-end coefficient $-\sigma_0^2/2$ is conserved at every vol-of-vol, unrenormalised (Theorem 6.4), and what stochastic intensity costs instead is an existence boundary, exact for the infinite-span surrogate (Theorem 6.5). In the *bilinear completion convention*, $(s_t, \mathbf{X}_t)$ is jointly Gaussian, the product $s_t \mathbf{X}_t$ is represented literally as a quadratic state, and the closure is studied under an explicit recovery normalization (§§6.5–6.7; Remark 6.13 records what the normalization leaves aside). There no-arbitrage locks

the generator and requires spread mean reversion to be slower than the target deformation decay (Theorem 6.12), and the same short-end coefficient is recovered on the $\eta = 0$ slice (Proposition 6.15). Existence of the finite-dimensional coefficient system at each truncation is decided by a determinant (Proposition 6.16); the short-end value at positive vol-of-vol is the open problem of §7.3. The affine convention answers the objection: in its uncorrelated form the short-end coefficient belongs to the maturity-local noise. Intensity scales it and adds its own terms away from the short end, but does not renormalise it, while the rest of the loading and its existence domain do change with the intensity. The bilinear convention makes visible what no-arbitrage demands of the spread itself.

6.1 The affine completion and its closure

The first convention makes the squared intensity $v_t = s_t^2$ a CIR factor, so that the deformation is conditionally Gaussian and the closure of §5 can be redone with a stochastic source. Statements are at general N because the $\mathbf{X}$ -sector closure is block-diagonal and the v -sector is scalar with additive source: the one-block calculation extends by summation.

Assumption 6.1 (Affine completion). Fix N and the data of Definition 5.1. The curve family is

$$f_t(\tau) = \varphi_t(\tau) + \psi(\tau) v_t + \mathbf{B}(\tau)^\top \mathbf{X}_t,$$

with φ_t deterministic, $\psi \in C^1$, $\psi(0) = 0$ (a normalisation; Remark 6.9), and dynamics

$$d\mathbf{X}_t = K\mathbf{X}_t dt + \sqrt{v_t} \Sigma d\mathbf{W}_t, \quad dv_t = \lambda(\bar{v} - v_t) dt + \eta\sqrt{v_t} dZ_t,$$

$\lambda, \bar{v} > 0$, $2\lambda\bar{v} \geq \eta^2$ (Feller), Σ the sine-only loading of Definition 5.1, and $d\langle W_{\sin}^n, Z \rangle_t = \rho_n dt$ with $\sum_n \rho_n^2 \leq 1$. Normalise $\bar{v} = 1$, so the frozen completion is recovered as $\eta \rightarrow 0$, $v_0 = 1$. This fixes the reference initialisation; consistency in Proposition 6.2 means that the drift identity holds for every initial (x, v) in a nonempty open subset of $\mathbb{R}^{2N} \times (0, \infty)$.

Proposition 6.2 (Closure; conditional locking). Write $b_n^s(\tau) = \sqrt{2/\theta^*} e^{-\tau/\kappa} \sin \omega_n \tau$, $I_n^s(\tau) = \int_0^\tau b_n^s$, and $\Psi(\tau) := \int_0^\tau \psi$. The model of Assumption 6.1 satisfies the drift condition (3), with the preceding every-initial-state consistency quantifier, if and only if

(i) $K = M^\top$;

(ii) ψ solves

$$\psi'(\tau) = -\lambda\psi(\tau) - q^N(\tau) - \eta^2\psi(\tau)\Psi(\tau) - \eta\sigma_0 \sum_{n \leq N} \rho_n [b_n^s(\tau)\Psi(\tau) + \psi(\tau)I_n^s(\tau)]; \quad (27)$$

(iii) $\dot{\varphi}_t(\tau) = \varphi'_t(\tau) - \lambda\bar{v}\psi(\tau)$.

For $\rho = 0$ the integrated form of (27) is the CIR-type Riccati

$$\Psi' = -\lambda\Psi - \frac{\eta^2}{2}\Psi^2 - Q^N, \quad Q^N(\tau) := \int_0^\tau q^N(s) ds. \quad (28)$$

Spectral locking survives conditionally: $\text{spec } K = \text{spec } M$ on the deformation sector, so Corollary 4.5's signature holds for the conditionally Gaussian factors, with one additional real eigenvalue $-\lambda$, which does not precess (Proposition 6.8).

Proof. Volatility matching gives $\sigma_t(\tau) = \sqrt{v_t} (\mathbf{B}(\tau)^\top \Sigma, \eta\psi(\tau))$, so the required drift (3) is $v_t[q^N + \eta^2\psi\Psi + \eta\sigma_0 \sum_n \rho_n(b_n^s\Psi + \psi I_n^s)]$ — linear in v_t with no constant part, consistent with the diffusion vanishing at $v = 0$. The model's drift is $\dot{\varphi}_t + \lambda(\bar{v} - v_t)\psi + \mathbf{B}^\top K \mathbf{X}_t$, while (2) demands $\varphi'_t + \psi'v_t + \mathbf{B}^\top M^\top \mathbf{X}_t + \alpha_t$. Matching $\mathbf{X}$ -linear terms forces $K = M^\top$ exactly as in Proposition 2.7 (consistency in the state, then Lemma 2.4), and v -linear terms give (27). Deterministic terms give (iii), whose solution is $\varphi_t(\tau) = \varphi_0(\tau + t) - \lambda\bar{v} \int_0^t \psi(\tau + t - s) ds$; the mean intensity feeds the loading through the deterministic sector, and no global sign of ψ is needed. Blocks are additive in both the $\mathbf{X}$ -matching and the source of (27), which reduces general N to $N = 1$. $\square$

Remark 6.3 (Consistency; strictness). (i) At $\eta = 0, \lambda = 0$: $v \equiv v_0$ and $\psi = -\int_0^\tau q^N$; the identity $\psi(\tau) - \psi(\tau + t) = \int_\tau^{\tau+t} q^N(s) ds$ (immediate from $\psi = -\int_0^\tau q^N$) shows the static loading $v_0\psi$ reproduces the transport (6) of the frozen completion exactly — the two constructions are the same model in different coordinates. (ii) At $\eta = 0, \lambda > 0$ (deterministic time-varying intensity, still inside Proposition 2.7's class), $\psi^N = -\int_0^\tau e^{-\lambda(\tau-s)} q^N(s) ds$ is quasi-exponential with exponent set

$$\{-\lambda\} \cup \{-\tfrac{1}{\kappa} \pm i\omega_n\} \cup \{-\tfrac{2}{\kappa}\} \cup \{-\tfrac{2}{\kappa} \pm 2i\omega_n\},$$

the beat lattice of Remark 2.11 plus one intensity mode: the lattice closure under beats — the “structural felicity” — is load-bearing here, keeping the strict realisation finite. (iii) For $\eta \neq 0$, equation (27) is nonlinear and the quasi-exponential closure argument in (ii) no longer applies. The curve family remains finite-dimensional and affine.

6.2 The short-end cost under stochastic intensity

With the affine closure in hand, the question is whether the short-end cost of Theorem 5.2 survives the stochastic intensity or is renormalised away by it. It survives, one derivative down.

Theorem 6.4 (The cost is conserved, demoted one derivative, and unrenormalised). *Let ψ^N solve (27) with source q^N , and ψ^∞ the solution with source $q^\infty = \frac{\sigma_0^2}{2} e^{-2\tau/\kappa}$, both with $\psi(0) = 0$ and the same parameters, with $\rho = 0$ for the surrogate comparison. Then:*

- (a) *Let τ_∞ be the maximal existence time of the surrogate Riccati flow driven by q^∞ , and let τ_N denote the (generally different) maximal time for the finite- N flow. For every $T < \min\{\theta^*, \tau_\infty\}$, one has $\tau_N > T$ for all sufficiently large N and*

$$\sup_{0 \leq \tau \leq T} |\psi^N(\tau) - \psi^\infty(\tau)| \longrightarrow 0.$$

Thus the θ^ -independent surrogate is the rigorous loading limit on the first span.*

- (b) *For every finite N and every admissible ρ , $\psi^N(0) = (\psi^N)'(0) = 0$. For the $\rho = 0$ surrogate, $(\psi^\infty)'(0^+) = -\frac{\sigma_0^2}{2}$ exactly for every $(\lambda, \eta, \bar{v})$. Vol-of-vol does not renormalise this right derivative (a statement at $\tau = 0^+$, on the existence interval, which always contains a neighbourhood of 0).*
- (c) *The first boundary layer survives one derivative down: uniformly for u in compact subsets of $[0, \infty)$,*

$$(\psi^N)' \left(\frac{u\theta^*}{N\pi} \right) \longrightarrow -\sigma_0^2 G(u).$$

In particular the limiting first-layer profile has a strict local minimum $-G(\pi)\sigma_0^2$ at $u = \pi$.

(d) At $\eta = 0$, $\psi^\infty(\tau) = -\frac{\sigma_0^2\kappa}{2} \frac{e^{-2\tau/\kappa} - e^{-\lambda\tau}}{\kappa\lambda - 2}$, degenerating to $-\frac{\sigma_0^2}{2} \tau e^{-2\tau/\kappa}$ at the resonance $\lambda = 2/\kappa$.

Proof. (b) is structural: at $\tau = 0^+$ every term of (27) except the source vanishes ($\psi(0) = \Psi(0) = b_n^s(0) = 0$), so $\psi'(0^+) = -q(0^+)$; $q^N(0) = 0$ for every N while $q^\infty(0^+) = \sigma_0^2/2$ (Theorem 5.2). (a) Take $\rho = 0$ and set $\delta = \psi^N - \psi^\infty$, $D = \Psi^N - \Psi^\infty$. Subtracting the two copies of (27) gives not a scalar equation for δ but the linear system

$$D' = \delta, \quad \delta' = -(\lambda + \eta^2 \Psi^N) \delta - \eta^2 \psi^\infty D - (q^N - q^\infty),$$

whose coefficients are bounded on any interval where the two flows exist. The Fourier representation in the proof of Theorem 5.2 gives a uniform bound for the partial sums on $[0, T]$ and pointwise convergence away from 0; dominated convergence therefore yields $q^N \rightarrow q^\infty$ in $L^1(0, T)$. Gronwall applied to $|D| + |\delta|$, together with standard ODE continuation, gives existence through T for all large N and the stated uniform convergence. This argument uses $T < \theta^*$ essentially. (c) In the layer $\psi^N = O(\tau)$ and $\Psi^N = O(\tau^2)$ uniformly on compact u -windows, so $(\psi^N)' = -q^N + o(1)$ and (25) applies with a sign flip. (d) Direct integration of the linear equation $\psi' = -\lambda\psi - q^\infty$. $\square$

Under the limit, $v_t q^\infty(\tau)$ is the maturity-white/ $\mathbf{X}$ -sector contribution to the HJM drift. When $\eta > 0$, the full drift also contains the intensity and cross-covariance terms displayed in (27): the source term alone is not the whole instantaneous convexity.

6.3 Existence, and what the surrogate costs

A Riccati loading can escape at finite maturity, and the affine completion is usable only where it does not. On an idealised infinite-span source the boundary is exact.

Theorem 6.5 (Existence boundary). *Take $\rho = 0$ and, as an idealised infinite-span surrogate, the source $q^\infty(\tau) = \frac{\sigma_0^2}{2} e^{-2\tau/\kappa}$ on all of $[0, \infty)$ — the interior limit of Theorem 5.2 extended beyond the span (Remark 6.6 records the fixed-span correction). The solution of (28) is global on $[0, \infty)$ if and only if*

$$\lambda \geq \eta\sigma_0\sqrt{\kappa/2}, \quad \text{equivalently} \quad \lambda^2 \geq 2\eta^2 Q_\infty, \quad Q_\infty := \int_0^\infty q^\infty(s) ds = \frac{\sigma_0^2\kappa}{4}.$$

Below the boundary, $\Psi \downarrow -\infty$ at a finite $\tau^(\lambda, \eta)$, and bond prices — exponential-affine with v -exponent $-\Psi(\tau)$ — are infinite for maturities beyond τ^* ; $\tau^* \uparrow \infty$ as $\lambda \uparrow \eta\sigma_0\sqrt{\kappa/2}$. The criterion is the CIR exponential-moment condition for $\mathbb{E}[\exp(Q_\infty \int_0^\cdot v_s ds)] < \infty$; see [4] for the moment criterion.*

Proof. $\Phi := -\Psi \geq 0$ satisfies $\Phi' = \frac{\eta^2}{2} \Phi^2 - \lambda\Phi + Q(\tau)$ with $Q(\tau) = \int_0^\tau q^\infty \uparrow Q_\infty$. If $\eta = 0$, this is a linear equation and is global. Assume henceforth $\eta > 0$. If $\lambda^2 \geq 2\eta^2 Q_\infty$, the autonomous majorant with source Q_∞ has fixed points $\Phi_\pm = [\lambda \pm \sqrt{\lambda^2 - 2\eta^2 Q_\infty}]/\eta^2$, and the trajectory from $\Phi(0) = 0 < \Phi_-$ is trapped below Φ_- : global existence. If $\lambda^2 < 2\eta^2 Q_\infty$, choose T_0 with $Q(\tau) \geq \tilde{Q} > \lambda^2/(2\eta^2)$ for $\tau \geq T_0$. The autonomous minorant with source $\tilde{Q}$ has no fixed points and a vector field bounded below by a positive constant plus the quadratic, so it blows up

in finite time, and comparison forces Φ to do the same. Finally, on the subcritical side the blow-up time is bounded below by that of the autonomous majorant with source Q_∞ , whose vector field has no real roots, so

$$\tau^* \geq \int_0^\infty \frac{d\phi}{\frac{\eta^2}{2}\phi^2 - \lambda\phi + Q_\infty} = \frac{\pi + 2 \arctan(\lambda/\sqrt{2\eta^2 Q_\infty - \lambda^2})}{\sqrt{2\eta^2 Q_\infty - \lambda^2}} \longrightarrow \infty \quad \text{as } \lambda \uparrow \eta\sigma_0\sqrt{\kappa/2},$$

which is the stated divergence of τ^* . $\square$

Remark 6.6 (The fixed-span completion; surrogate accuracy). Theorem 5.2 identifies q^∞ only on $(0, \theta^*)$. The quasi-exponential continuation of the fixed-span completion differs beyond the span: on $(\theta^*, 2\theta^*)$,

$$q_{\text{span}}^\infty(\tau) = \frac{\sigma_0^2}{2} e^{-2\tau/\kappa} - \sigma_0^2 e^{-2\theta^*/\kappa}$$

— negative on this second span. This follows by inserting $\tau = \theta^* + u$, $0 < u < \theta^*$, into the periodic Fourier sums in the proof of Theorem 5.2: the first-harmonic sum changes sign, whereas the even-harmonic sum is unchanged. Thus its left and right limits at θ^* are respectively $\frac{\sigma_0^2}{2} e^{-2\theta^*/\kappa}$ and $-\frac{\sigma_0^2}{2} e^{-2\theta^*/\kappa}$: the full two-sided jump has size $\sigma_0^2 e^{-2\theta^*/\kappa}$, although every finite sum takes the midpoint value $q^N(\theta^*) = 0$.

For completeness, the fixed-span continuation is well-defined on the whole half-line. In the Fourier representation used in Theorem 5.2, the partial sine sums with coefficients $\omega_n/(\kappa^{-2} + \omega_n^2)$ are uniformly bounded, while the remaining coefficient sequences are absolutely summable. Consequently, for a constant C independent of N ,

$$|q^N(\tau)| \leq C(e^{-\tau/\kappa} + e^{-2\tau/\kappa}), \quad \tau \geq 0.$$

The Fourier series converges pointwise away from the lattice $\theta^*\mathbb{N}$ (and to its midpoint values on that null set); call the a.e. limit q_{span}^∞ . Dominated convergence on $[0, \infty)$ then gives

$$Q^{\text{span}}(\tau) := \int_0^\tau q_{\text{span}}^\infty(s) ds, \quad Q^{\text{span}}(\tau) \longrightarrow \lim_{N \rightarrow \infty} \int_0^\tau q^N(s) ds.$$

Termwise evaluation of the absolutely convergent terminal series yields the true terminal convexity mass

$$Q_\infty^{\text{span}} := \lim_{N \rightarrow \infty} Q^N(\infty) = \frac{\sigma_0^2 \kappa}{4} [\coth y - y \operatorname{csch}^2 y], \quad y = \frac{\theta^*}{\kappa},$$

strictly below the surrogate's $Q_\infty = \sigma_0^2 \kappa/4$. The preceding domination proves both $Q^{\text{span}}(\tau) \rightarrow Q_\infty^{\text{span}}$ and $\sup_{\tau \geq 0} |Q^{\text{span}}(\tau)| < \infty$. This terminal mass does *not* give an if-and-only-if existence criterion because Q^{span} need not be monotone and can overshoot its limit. Writing $\overline{Q}^{\text{span}} := \sup_{\tau \geq 0} Q^{\text{span}}(\tau) < \infty$, comparison for $\Phi = -\Psi$ gives the rigorous bracket

$$\lambda^2 \geq 2\eta^2 \overline{Q}^{\text{span}} \implies \text{global existence}, \quad \text{global existence} \implies \lambda^2 \geq 2\eta^2 Q_\infty^{\text{span}}.$$

The first implication uses the autonomous majorant with source $\overline{Q}^{\text{span}}$; the second follows because if the terminal discriminant is negative, the source is eventually bounded below by a supercritical constant. Finally,

$$\frac{Q_\infty^{\text{span}}}{Q_\infty} = 1 + (2 - 4y)e^{-2y} + O(ye^{-4y}),$$

so the relative correction is $O(ye^{-2y})$, about 10^{-5} at the reference $y = 7.5$.

Remark 6.7 (DIR companion). Conditionally Gaussian rates are unbounded below, so bond prices are exponential moments of integrated intensity; Theorem 6.5 says the *total convexity mass* $\sigma_0^2 \kappa / 4$ must be an admissible such moment. This is the not-innocent-limit phenomenon of §5.3 resurfacing in the vol-of-vol direction: the affine extension prices maturity-white risk globally only while mean reversion dominates $\eta \sigma_0 \sqrt{\kappa} / 2$. The “not innocent” caution of [13] acquires a second exact instance.

6.4 Where complex spectrum may not live

The deformation sector carries rotation freely. The intensity sector cannot, and the reason is positivity rather than anything spectral.

Proposition 6.8 (Precession constraints in the intensity sector). *Suppose more generally that m intensity factors $v \in \mathbb{R}_+^m$ drive the deformation covariance, with affine dynamics $dv = (a + \mathcal{K}v) dt + \text{diag}(\eta_i \sqrt{v_i}) dZ$. Invariance of the cone $\mathbb{R}_+^m$ (admissibility, [14, 18]) requires the drift to point inward on each face: $a \geq 0$ and $\mathcal{K}$ Metzler (nonnegative off-diagonal). For $m \leq 2$ the spectrum of a Metzler matrix is real. For $m \geq 3$ complex pairs occur (cyclic feedback) but are confined to a Perron sector: rotation is bounded by decay relative to the Perron root (Dmitriev–Dynkin–Karpelevich-type sector theorems; [5]). In particular the model’s scalar spread ($m = 1$) cannot precess, and no two-factor intensity can. Higher-dimensional positive cones may carry complex pairs, but only under the Perron-sector restrictions just described; in particular they do not admit an unconstrained isolated two-dimensional rotation block.*

Proof. Admissibility. At a face point $v \geq 0$ with $v_i = 0$, the i -th diffusion entry $\eta_i \sqrt{v_i}$ vanishes, so invariance of the cone requires the drift to satisfy $a_i + \sum_{j \neq i} \mathcal{K}_{ij} v_j \geq 0$ there; $v = 0$ gives $a_i \geq 0$, and $v = T e_j$ with $T \rightarrow \infty$ gives $\mathcal{K}_{ij} \geq 0$ for $j \neq i$. (With the Feller strengthening on each face these conditions are also sufficient; [18, 14].)

$m \leq 2$. $m = 1$ is trivial. For $m = 2$ the characteristic polynomial of $\mathcal{K}$ has discriminant $(\mathcal{K}_{11} - \mathcal{K}_{22})^2 + 4\mathcal{K}_{12}\mathcal{K}_{21} \geq 0$ because the off-diagonal entries are nonnegative: the spectrum is real. (Equivalently, read on a rotation cell: the boundary drift $\dot{v}_2|_{v_2=0} = -\omega v_1$ must be nonnegative, forcing $\omega = 0$.)

Isolated blocks. Every principal submatrix of a Metzler matrix is Metzler; a decoupled 2×2 intensity block would be 2×2 Metzler, hence has real spectrum. No admissible intensity sector, in any dimension, contains an isolated rotation cell.

Perron sector for $m \geq 3$. Set $c_0 := \max_i (-\mathcal{K}_{ii})^+$, so $P := \mathcal{K} + c_0 \mathbf{1} \geq 0$ entrywise. By Perron–Frobenius [5], $\rho(P) \in \text{spec } P$, so the spectral abscissa $\alpha := \max \text{Re spec } \mathcal{K}$ is attained by a real eigenvalue and $\rho(P) = \alpha + c_0$. Any eigenvalue $\mu = \xi + i\zeta$ of $\mathcal{K}$ has $|\mu + c_0| \leq \rho(P)$, i.e.

$$\zeta^2 \leq (\alpha - \xi)(\alpha + \xi + 2c_0) :$$

rotation is bounded by the decay gap to the Perron root, at the scale set by the diagonal. Complex pairs do occur — the cyclic 3×3 permutation matrix has spectrum $\{1, e^{\pm 2\pi i/3}\}$ and saturates the bound — but an unconstrained rotation pair, $\xi = \alpha$ with $\zeta \neq 0$, is excluded, since there $\zeta^2 \leq 0$. Sharper sector theorems (Dmitriev–Dynkin–Karpelevich) are in [5]. In the conditionally Gaussian construction the rotation cells sit in the $\mathbf{X}$ -fibre, unconstrained by positivity; only the intensity carries the cone restriction, and for the model that intensity is scalar. $\square$

Remark 6.9 (Gauge freedom; what remains open). $\psi(0) = 0$ is a normalization, not a necessity: changing it changes the short rate's direct intensity loading and hence the Riccati source. Likewise, adding a vector $g_N v_t$ to the $\mathbf{X}$ -drift redistributes terms between the scalar and factor sectors. Any assertion uniform in N must impose a structured-gauge condition such as $\sup_N |\mathbf{B}_N(0)^\top g_N| < \infty$; a uniform ℓ^2 bound on g_N alone does not imply this, because $\mathbf{B}_N(0)$ has N nonzero cosine coordinates. The completion freedom of Remark 5.3(iii) therefore persists. The gauge condition stated here is what is needed for conclusions that are uniform over finite truncations.

That closes the affine convention. The completion's closure exists with a stochastic source (Proposition 6.2). The short-end coefficient is conserved and unrenormalised at every vol-of-vol (Theorem 6.4). Its price is an existence boundary, exact for the infinite-span surrogate and bracketed for the fixed span (Theorem 6.5, Remark 6.6). And rotation stays confined to the deformation sector (Proposition 6.8).

The affine convention is, however, a surrogate: it carries the spread through its square. The model's own algebra multiplies the deformation by s_t itself, so $s_t \mathbf{B}(\tau)^\top \mathbf{X}_t$ is bilinear when $(s_t, \mathbf{X}_t)$ is jointly Gaussian, and the literal reading lies in the quadratic-Gaussian class [2, 34]. Within polynomial term-structure families, degree two is the last nontrivial case under the hypotheses of the maximal-degree theorem [20], so the bilinear family is the natural terminal member of that class. The construction below fixes a recovery normalization and studies the finite-dimensional HJM closure it induces; what no-arbitrage asks of the spread itself becomes visible only here.

6.5 The bilinear channel: quadratic-Gaussian family and normalization

Here is the family, together with the normalization under which its closure is studied.

Assumption 6.10 (Quadratic-Gaussian family and recovery normalization). State $Z_t = (s_t, \mathbf{X}_t) \in \mathbb{R}^{1+2N}$ with

$$dZ_t = (\mathcal{A}Z_t + a) dt + \Lambda d\mathbf{W}_t, \quad \mathcal{A} = \begin{pmatrix} -\beta & 0 \\ 0 & K \end{pmatrix}, \quad a = \begin{pmatrix} \beta \bar{s} \\ \mathbf{0} \end{pmatrix},$$

$\beta, \bar{s} > 0$, $V := \Lambda \Lambda^\top$ with spread variance η^2 , the sine-only deformation loading C of Definition 5.1, and spread-deformation correlation ρ (set to 0 in the displayed specialization). Curve family

$$f_t(\tau) = \varphi_t(\tau) + \mathbf{b}(\tau)^\top Z_t + Z_t^\top \mathbf{G}(\tau) Z_t, \quad \mathbf{b} = \begin{pmatrix} Y \\ \mathbf{b}_X \end{pmatrix}, \quad \mathbf{G} = \begin{pmatrix} g & \frac{1}{2} \gamma^\top \\ \frac{1}{2} \gamma & H \end{pmatrix},$$

where $\mathbf{G}(\tau) \in \mathbb{S}^{1+2N}$ and $H(\tau) \in \mathbb{S}^{2N}$ are symmetric. The symmetry is part of the parametrisation, not a restriction on the model: adding a skew $H_a(\tau)$ leaves the curve f_t unchanged while altering the coefficient system, so only the symmetric part is identified. The family is indexed by η with $\beta, \bar{s}, K$ and the deformation loading C held fixed, which is what allows a normalization imposed at $\eta = 0$ to constrain the family. The *recovery normalization* is a model choice: for every η in the family set $\gamma_\eta(0) = \mathbf{B}(0)$, together with $g_\eta(0) = 0$, $H_\eta(0) = \mathbf{0}$, and $\mathbf{b}_{X,\eta}(0) = \mathbf{0}$; at $\eta = 0$ require in addition that the bilinear loading reproduce the completed deformation loading, $\gamma_0(\tau) \equiv \mathbf{B}(\tau)$. The normalization is not a consequence of no-arbitrage.

Without it, the $\eta = 0$ bilinear equation admits, for every K , the loading $\gamma(\tau) = e^{(K^\top - \beta \mathbf{1})\tau} \mathbf{B}(0)$, and nothing locks.

On every maturity interval considered below, $\mathbf{b}$ and $\mathbf{G}$ are C^1 in maturity, φ is jointly C^1 on the corresponding maturity–time triangles, and the integrability hypotheses of Lemma 2.1 hold. The matching equations below are identities between coefficients, maturity by maturity.

Proposition 6.11 (QG–HJM matching). *Require (3) for every initial state in a nonempty open subset of $\mathbb{R}^{1+2N}$, equivalently as a polynomial identity in the current state. The model of Assumption 6.10 is then consistent if and only if, with $\mathbf{G}_I(\tau) := \int_0^\tau \mathbf{G}$ and $\mathbf{b}_I := \int_0^\tau \mathbf{b}$,*

$$\mathbf{G}' = \mathcal{A}^\top \mathbf{G} + \mathbf{G} \mathcal{A} - 2(\mathbf{G} V \mathbf{G}_I + \mathbf{G}_I V \mathbf{G}), \quad \mathbf{b}' = \mathcal{A}^\top \mathbf{b} + 2\mathbf{G} a - 2(\mathbf{G}_I V \mathbf{b} + \mathbf{G} V \mathbf{b}_I),$$

and $\dot{\varphi}_t = \varphi'_t + \mathbf{b}^\top V \mathbf{b}_I - a^\top \mathbf{b} - \text{tr}(\mathbf{G} V)$.

Proof. For fixed τ , Itô's formula gives the family's drift, by polynomial degree in z , as

$$\dot{\varphi}_t + \mathbf{b}^\top a + \text{tr}(\mathbf{G} V) + (\mathcal{A}^\top \mathbf{b} + 2\mathbf{G} a)^\top z + z^\top (\mathcal{A}^\top \mathbf{G} + \mathbf{G} \mathcal{A}) z,$$

and its Brownian loading is $(\mathbf{b} + 2\mathbf{G} z)^\top \Lambda$. The HJM covariance term is

$$(\mathbf{b} + 2\mathbf{G} z)^\top V (\mathbf{b}_I + 2\mathbf{G}_I z).$$

Its constant, linear, and symmetric quadratic coefficients are respectively

$$\mathbf{b}^\top V \mathbf{b}_I, \quad 2(\mathbf{G}_I V \mathbf{b} + \mathbf{G} V \mathbf{b}_I), \quad 2(\mathbf{G} V \mathbf{G}_I + \mathbf{G}_I V \mathbf{G}).$$

Equating these coefficients with those of $\partial_\tau f_t$ proves the displayed system and the scalar transport equation. Conversely, substitution makes the polynomial drift identity hold, proving sufficiency. $\square$

That is the matching system; the next subsection reads off what the recovery normalization locks in it.

6.6 What the normalization forces

Matching is now a matter of reading the block system, and three things fall out of it: the generator is locked, the spread must be slowly varying, and the family is not closed under no-arbitrage.

Theorem 6.12 (Structure forced by the recovery normalization). *Impose the recovery normalization. Then:*

- (a) *(Conditional gauge locking.) Given the recovery normalization, the bilinear block forces $K = M^\top + \beta \mathbf{1}$, uniquely; no-arbitrage alone does not lock (Assumption 6.10). Only the products are observable: the homogeneous bilinear $s_t \mathbf{X}_t$ block has drift coefficient $M^\top$. When $\bar{s} \neq 0$, the full product drift also contains the affine source $\beta \bar{s} \mathbf{X}_t$; correspondingly the augmented observable system retains the spectrum of K . The split of homogeneous decay between s and $\mathbf{X}$ is invisible in the product: $(s, \mathbf{X}) \mapsto (cs, \mathbf{X}/c)$ fixes $s_t \mathbf{X}_t$ for every $c \neq 0$, and it extends to a gauge of the whole curve family only when the linear and quadratic loadings and the parameters are transformed to match. Spectral locking (Corollary 4.5) holds on the homogeneous bilinear block.*

- (b) (The spread must be slowly varying.) $\text{spec } K = \beta - \frac{1}{\kappa} \pm i\omega_n$: the state is stationary iff $\beta < 1/\kappa$ — the spread must mean-revert more slowly than the deformation decays.
- (c) (Generated blocks.) The bilinear family is not closed: the drift condition generates exactly three completion blocks — an s^2 -loading g sourced by the convexity, a matrix $\mathbf{X}\mathbf{X}^\top$ -loading H sourced by the spread's vol-of-vol, and a linear $\mathbf{X}$ -loading $\mathbf{b}_X$ sourced by the mean spread. At $\rho = 0$ the closed system is

$$\begin{aligned}
g' &= -2\beta g - \gamma^\top C \gamma_I - 4\eta^2 g g_I, \\
\gamma' &= M\gamma - 2\eta^2(g\gamma_I + g_I\gamma) - 2(HC\gamma_I + H_IC\gamma), \\
H' &= K^\top H + HK - 2(HCH_I + H_ICH) - \frac{\eta^2}{2}(\gamma\gamma_I^\top + \gamma_I\gamma^\top), \\
Y' &= -\beta Y + 2\beta\bar{s}g - 2\eta^2(Yg_I + Y_Ig) - (\gamma_I^\top C \mathbf{b}_X + \gamma^\top C \mathbf{b}_{X,I}), \\
\mathbf{b}'_X &= K^\top \mathbf{b}_X + \beta\bar{s}\gamma - \eta^2(Y\gamma_I + Y_I\gamma) - 2(HC \mathbf{b}_{X,I} + H_IC \mathbf{b}_X),
\end{aligned} \tag{29}$$

Here subscript I denotes $\int_0^\tau$. At $\eta = 0$, $H \equiv 0$ and $\gamma = \mathbf{B}$ exactly, while

$$g(\tau) = - \int_0^\tau e^{-2\beta(\tau-u)} q^N(u) du.$$

This is quasi-exponential, with $\{-2\beta\}$ adjoined to the beat lattice of Remark 6.3.

Proof. (a) The $s\mathbf{X}$ -block of the $\mathbf{G}$ -equation with the recovery normalization $\gamma = \mathbf{B}$ at $\eta = 0$ reads $\mathbf{B}' = (K^\top - \beta\mathbf{1})\mathbf{B}$, i.e. $(K^\top - \beta\mathbf{1} - M)\mathbf{B}(\tau) = 0$ for all τ : the constant matrix annihilates $\text{span}\{\mathbf{B}(\tau) : \tau \geq 0\}$, which is all of $\mathbb{R}^{2N}$ by controllability (Lemma 2.4), so $K = M^\top + \beta\mathbf{1}$, uniquely. The homogeneous bilinear coefficient in $d(s\mathbf{X})$ is $-\beta\mathbf{1} + K = M^\top$; when $\bar{s} \neq 0$ the same Itô calculation also gives the affine term $\beta\bar{s}\mathbf{X} dt$. (b) Immediate from (a). (c) Substitute the block forms of $\mathbf{G}$, $\mathcal{A}$, a , V into the matrix system of Proposition 6.11 and read off the (s, s) -, $(s, \mathbf{X})$ - and $(\mathbf{X}, \mathbf{X})$ -components: this yields (29), and the three sources are the ss -, $\mathbf{X}\mathbf{X}$ -, and s -linear components of $-2(\mathbf{G}V\mathbf{G}_I + \mathbf{G}_IV\mathbf{G})$ and $2\mathbf{G}a$. At $\eta = 0$ the H -source vanishes and $H(0) = 0$ gives $H \equiv 0$, whence $\gamma = \mathbf{B}$ exactly and the g -line integrates in closed form. $\square$

Remark 6.13 (Endpoint scope). The model of [13] imposes $Y(0) = 0$, $Y(\theta^*) = 1$, and Dirichlet conditions on the deformation. Those conditions are not imposed by the recovery system. In particular, $\gamma(0) = \mathbf{B}(0)$ contains cosine coordinates, whereas the model's deformation vanishes at the short endpoint. This is therefore a completion of the selected bilinear algebra under a recovery normalization, not a literal recovery of the model's endpoint conditions.

Remark 6.14 (Linear-quadratic relation). By Itô, $d(s^2) = [\eta^2 + 2\beta\bar{s}s - 2\beta s^2]dt + 2\eta s dZ$; setting $d\tilde{Z} := \text{sgn}(s_t)dZ$, a Brownian motion by Lévy's characterisation, turns the diffusion into $2\eta\sqrt{\bar{v}}d\tilde{Z}$. In the central case $\bar{s} = 0$ — outside Assumption 6.10's $\bar{s} > 0$, and taken here only as the comparison point — the squared spread is then CIR with $(\lambda, \bar{v}, \eta_v) = (2\beta, \frac{\eta^2}{2\beta}, 2\eta)$, at the Feller-violating corner, as it must be (an OU process crosses zero). The loading comparison below is unaffected by $\bar{v}$, which enters only the deterministic transport. Under this map the g -line of (29) with γ frozen to $\mathbf{B}$ is *identical* to the ψ -equation (27). So §6.1 is the central, level-blind case (cf. the linear-quadratic equivalence of [10]). The additional QG content is exactly two channels absent from §6.1: the γ -flow, i.e. curve volatility from $X ds$ — the deformation *level* times the spread's vol-of-vol — and the matrix sector H .

Proposition 6.15 (Finite-truncation short-end identities). *For every finite N and every parameter set on the local existence interval, $g^N(0) = (g^N)'(0) = 0$. In the exactly solvable case $\eta = 0$ one has $H \equiv 0$, $\gamma = \mathbf{B}$, and, on the first span after passing to the interior completion limit,*

$$g^\infty(\tau) = -\frac{\sigma_0^2 \kappa}{2} \frac{e^{-2\tau/\kappa} - e^{-2\beta\tau}}{2\beta\kappa - 2},$$

with the resonant value $g^\infty(\tau) = -\frac{\sigma_0^2}{2}\tau e^{-2\tau/\kappa}$ at $\beta = 1/\kappa$. Consequently $(g^\infty)'(0^+) = -\sigma_0^2/2$, and on fixed first-layer windows

$$(g^N)' \left(\frac{u\theta^*}{N\pi} \right) \longrightarrow -\sigma_0^2 G(u).$$

Proof. At $\tau = 0$ all integrated loadings in the g -equation vanish, giving the finite- N identities. When $\eta = 0$, the system reduces to $g' + 2\beta g + q^N = 0$ with $g(0) = 0$. The displayed formula follows from Theorem 6.4(d) under $(\lambda, \eta_v) = (2\beta, 0)$; the first-layer limit follows exactly as in Theorem 6.4(c). $\square$

Both identities hold at every finite N . Whether the loading that carries them exists at every maturity is the remaining question.

6.7 Finite-truncation existence

What remains is whether the bilinear loading exists at every maturity. At fixed N that question has an exact answer, through the linearisation of the loading Riccati.

Proposition 6.16 (Finite-dimensional QG existence domain). *For every fixed finite N , the integrated loading solves the autonomous matrix Riccati equation*

$$\mathbf{G}'_I = \mathbf{G}(0) + \mathcal{A}^\top \mathbf{G}_I + \mathbf{G}_I \mathcal{A} - 2\mathbf{G}_I V \mathbf{G}_I, \quad \mathbf{G}_I(0) = 0.$$

If

$$\begin{pmatrix} U \\ W \end{pmatrix}' = \begin{pmatrix} -\mathcal{A} & 2V \\ \mathbf{G}(0) & \mathcal{A}^\top \end{pmatrix} \begin{pmatrix} U \\ W \end{pmatrix}, \quad U(0) = \mathbf{1}, \quad W(0) = 0,$$

then $\mathbf{G}_I = WU^{-1}$ until the first positive zero of $\det U$. Thus the finite- N flow is global if and only if $\det U(\tau) \neq 0$ for every $\tau > 0$. On this maximal interval, $\mathbf{G} = \mathbf{G}'_I$ is finite. For each prescribed value of the otherwise unconstrained scalar coefficient $Y(0)$ (with $\mathbf{b}_X(0) = \mathbf{0}$ as in Assumption 6.10) and each admissible initial deterministic curve φ_0 , the remaining loading equations for $\mathbf{b}$ and φ have unique solutions. Thus the same determinant gives the maximal maturity interval of the full finite-dimensional HJM coefficient system, subject to those initial data.

Proof. Integrating the first equation of Proposition 6.11 and using $\mathbf{G}_I(0) = 0$ gives the displayed autonomous Riccati equation. Wherever U is invertible, differentiation of WU^{-1} and substitution of the two linear equations give

$$(WU^{-1})' = \mathbf{G}(0) + \mathcal{A}^\top (WU^{-1}) + (WU^{-1})\mathcal{A} - 2(WU^{-1})V(WU^{-1}).$$

Uniqueness for the Riccati initial-value problem therefore yields $\mathbf{G}_I = WU^{-1}$. Conversely, if $\mathbf{G}_I$ exists on $[0, T]$, define U there by the linear equation $U' = (-\mathcal{A} + 2V\mathbf{G}_I)U$, $U(0) = \mathbf{1}$, and put $W := \mathbf{G}_I U$; the pair solves the displayed Hamiltonian system, and Liouville's formula

$$\det U(\tau) = \exp \int_0^\tau \text{tr}(-\mathcal{A} + 2V\mathbf{G}_I(u)) du$$

shows $\det U$ never vanishes on $[0, T]$. Hence the maximal interval ends precisely when the graph chart loses invertibility, i.e. at the first positive zero of $\det U$.

Once $\mathbf{G}_I$ is finite, the Riccati equation makes $\mathbf{G} = \mathbf{G}'_I$ continuous. Adjoining $\mathbf{b}'_I = \mathbf{b}$ rewrites the second matching equation as the inhomogeneous linear system

$$\frac{d}{d\tau} \begin{pmatrix} \mathbf{b}_I \\ \mathbf{b} \end{pmatrix} = \begin{pmatrix} 0 & \mathbf{1} \\ -2\mathbf{G}V & \mathcal{A}^\top - 2\mathbf{G}_I V \end{pmatrix} \begin{pmatrix} \mathbf{b}_I \\ \mathbf{b} \end{pmatrix} + \begin{pmatrix} 0 \\ 2\mathbf{G}a \end{pmatrix}.$$

Its coefficients are continuous on every compact subinterval of the Riccati existence interval, so, for $\mathbf{b}_I(0) = 0$ and the prescribed $\mathbf{b}(0)$, the loadings $\mathbf{b}_I$ and $\mathbf{b}$ extend uniquely throughout that interval. With

$$r(\tau) := \mathbf{b}^\top V \mathbf{b}_I - a^\top \mathbf{b} - \text{tr}(\mathbf{G}V),$$

the scalar equation is $\dot{\varphi}_t = \varphi'_t + r$. For prescribed φ_0 its characteristic solution is

$$\varphi_t(\tau) = \varphi_0(\tau + t) + \int_0^t r(\tau + t - s) ds,$$

wherever $\tau + t$ remains in the same interval. Hence no other coefficient can shorten the maximal maturity interval on the corresponding maturity–time triangle. $\square$

A zero of $\det U$ is the maturity at which the exponential–quadratic loading system ceases to exist, and with it the family's bond of that maturity: an existence boundary of the model, not an arbitrage strategy or an empirical estimate. For the $\rho = 0$ minimal recovery-gauge system, [42] proves that on each compact $[\beta_-, \beta_+] \subset (0, \kappa^{-1})$ there is an explicit $\eta_0 > 0$, independent of N , for which every finite truncation is global whenever $0 \leq \eta \leq \eta_0$.

The bilinear calculation is the finite-dimensional HJM closure of the selected spread–deformation algebra under the recovery normalization: it identifies the observable product drift and the polynomial blocks no-arbitrage forces, and Remark 6.13 records what the normalization leaves aside. Proposition 6.16 decides existence exactly at every fixed truncation; certifying the absence of conjugate points on the whole half-line at positive vol-of-vol is the global problem taken up in [42].

Within that scope the affine calculation answers the objection with which the section opened. In the uncorrelated affine convention, the maturity-white/ $\mathbf{X}$ -sector short-end coefficient $-\sigma_0^2/2$ is conserved at every vol-of-vol. Stochastic intensity nevertheless changes the remaining loading and can impose the existence boundary of Theorem 6.5. In the bilinear completion convention, no-arbitrage locks the homogeneous product generator, requires the stated stationarity inequality, and generates the additional completion blocks; the same coefficient is proved only on the exactly solvable $\eta = 0$ slice. Thus the invariant is the short-end coefficient of the $\mathbf{X}$ -sector source, not the full stochastic-intensity drift: the intensity scales that coefficient and adds its own terms away from the short end, but does not renormalise it. What the construction establishes, and where it stops, is collected in the conclusion.

7 Conclusion

The introduction separated two questions: whether a finite truncation of the cylindrical deformation, at fixed spread intensity, is consistent with pricing, and what that truncation preserves and what depends on its maturity resolution. It posed them for a source specification whose link to the pricing measure had been left open. This section answers them in the order the mathematics settles them, gives the financial reading of the answers, and states what remains open.

7.1 What has been shown

One observation carries the paper. The deformation's sine modes are the imaginary parts of complex exponentials, so the smallest shift-invariant family containing them is not their span but the doubled sine-cosine envelope (Proposition 4.1), and the coordinates the envelope adjoins carry no noise. Read on the envelope, the model's open link to $\mathbb{Q}$ becomes a decidable question with a located answer. Whether its factor drift is arbitrage-free is the polynomial divisibility $m_M \mid p - z$ (Theorem 4.3), with the convexity HJM adds carried by the deterministic sector as in Proposition 2.7. Here the model's symbol $p(z) = z + \kappa z^2/2$ fixes only $z = 0$, and its failure sits in the silent coordinates, beyond the reach of any change of measure.

Everything the paper says about pricing consistency follows from that location. Deterministic roll-down — the rotation ω_n acting on the cosine partners — lies outside the volatility span, so no finite sine truncation admits an equivalent local martingale measure with the model's own volatility. The discrepancy is realised by an explicit arbitrage in $N + 1$ bonds on an open dense set of maturity panels (Theorem 4.6), and its stationary size grows like $\sqrt{N}$ rather than fading under refinement (Proposition 4.8). In the infinite model the shift that would remove it has infinite Cameron–Martin energy (Proposition 4.10). Giving the silent coordinates noise of their own (§4.7) changes the finite verdict from an impossibility to a price: each truncation then admits an equivalent measure, at a Sharpe cost growing like $N^{3/2}$. The infinite verdict stands, since the infinite laws are singular (Proposition 4.11). That all of this is proved on finite-dimensional invariant subspaces is not a convenience: under cylindrical noise forward rates are not semimartingales (Proposition 3.3), so those subspaces are where this paper makes the question rigorous.

The same coordinates are where the repair goes. Adjoin the roll-down states, keep the sine-driven noise, correct only the drift, and every truncation becomes an ordinary $2N$ -factor Gaussian HJM model (§5). What the repair adds is the roll-down the model's span could not carry and the convexity HJM then requires. The convexity is where the cost lies, and the cost is exact: its refinement limit is $q^\infty(\tau) = \frac{\sigma_0^2}{2} e^{-2\tau/\kappa}$, a short-end jump of $\sigma_0^2/2$ reached through a boundary layer of width θ^*/N with the universal profile G (Theorem 5.2). Restoring the stochastic spread leaves this coefficient in place in the uncorrelated affine convention: $-\sigma_0^2/2$ is conserved at every vol-of-vol, while the rest of the loading changes and an existence boundary appears, exact for the infinite-span surrogate (Theorems 6.4 and 6.5). The short-end cost is therefore a property of the maturity-local noise, not of the frozen intensity. In the bilinear completion convention, no-arbitrage locks the homogeneous product generator and recovers the coefficient on the exactly solvable $\eta = 0$ slice (Theorem 6.12, Proposition 6.15).

Refinement answers the finite-implementation question first at the level of the coefficients entering bond returns. The convexity density converges locally uniformly on compact subsets of

$(0, \theta^*)$, and, after natural zero-padding, the bond-return volatility vectors converge in ℓ^2 ; every finite truncation prices its own bonds consistently. The short-end nonuniformity is resolved on the scale θ^*/N by the profile of Theorem 5.2, while Remark 6.6 records the separate far-end discontinuity. These are coefficient limits; the infinite traded-bond market itself is the first open problem of §7.3.

At the level of the forward curve the completions do not converge. Within the affine realisation class of Assumption 2.6, spectral locking replaces the model's $\lambda_n \asymp n^2$ by the uniform decay $1/\kappa$ (Corollary 4.5). That n^2 decay came from the maturity diffusion $\frac{\kappa}{2}\partial_\theta^2$ and made the cylindrical field function-valued (Propositions 3.2 and 5.4). The pricing counterpart established here is therefore a family of finite-factor bond models indexed by resolution, with exact limits for specified bond-level coefficients but no curve-valued refinement limit. Resolution is therefore where the economics enters, as §5.2 says: factor count is a modelling choice with an exact convexity attached, not a numerical grid parameter. Within this refinement and its minimal completion, a curve-valued limit can be obtained only by giving up maturity-local noise or no-arbitrage at every resolution; §5.2 shows what the first of these buys. The long end stands apart from all of this: the Dybvig–Ingersoll–Ross constraint is neither caused nor cured by complexification (Proposition 5.5). What the limit means for short-maturity volatility, and the witness for pricing, is the subject of the next subsection.

7.2 Economic interpretation

The completion prices its own finite-dimensional volatility specification; it does not assign a preference-free premium to the maturity-white component of the historical model of [13]. Once the sine-only volatility and the minimal sine–cosine envelope have been fixed, however, the HJM drift is no longer optional. Its first-span interior limit satisfies

$$q^\infty(\tau) = \frac{1}{2} \frac{d}{d\tau} \|\Sigma(\tau)\|^2 = \frac{\sigma_0^2}{2} e^{-2\tau/\kappa},$$

and hence $\|\Sigma(\tau)\|^2 \sim \sigma_0^2 \tau$ as $\tau \downarrow 0$. This is the model-internal economic meaning of the kink: short-maturity bond volatility scales as $\sigma_0 \sqrt{\tau}$ rather than linearly in τ , as it would for a smooth finite-factor volatility loading. The order of limits is $\lim_{\tau \downarrow 0} \lim_{N \rightarrow \infty}$ and the two do not commute: at fixed N , $I_n^s(\tau) = \sqrt{2/\theta^*} \omega_n \tau^2/2 + O(\tau^3)$, so $\|\Sigma(\tau)\| = O(\tau^2)$. The $\sqrt{\tau}$ law is a property of the refinement limit, not of any single truncation.

The boundary layer also records nonuniformity of implementation. At every finite N the drift density vanishes at the endpoint, while at each fixed interior maturity it approaches the nonzero limit above. The theorem therefore quantifies the local convexity cost of this chosen completion and the maturity scale on which successive truncations disagree.

The finite-bond witness gives the pricing-consistency problem a separate, security-level meaning. The original N -sine truncation does not merely fail a factor-coordinate identity: for generic maturities, $N + 1$ risky bonds contain a locally riskless self-financing direction whose discounted drift has a known sign after observing the current state. Likewise, $\mathfrak{d}_N^2$ is a normalized factor-quotient diagnostic: its linear growth shows that the unreachable drift does not become negligible per retained mode, and Remark 4.9 says what it does and does not measure. Under the isotropic loading of §4.7, by contrast, the invariant is a market-price-of-risk energy. There $\sqrt{\mathcal{E}_N}$ is the stationary root-mean-square of the maximal instantaneous Sharpe ratio of the N th truncated market, and its limit $\sqrt{\mathcal{E}}$ is finite exactly when the block defect energy $\sum_n |\delta_n|^2/k_n$ is (Proposition 4.11).

7.3 Limitations and open problems

The claims above stop at the following boundaries, each stated so that what would remove it is visible; see also Remarks 2.10, 4.4, 4.9, 5.3(iii), 6.9 and 6.13.

1. **Boundary conditions and domain.** The comparison extends the eigenmodes from $[0, \theta^*]$ to $\mathbb{R}_+$ by their quasi-exponential expressions; the Dirichlet conditions are respected only in the sense of Proposition 4.1, which shows precisely that they are incompatible with shift-invariance. The completion abandons them (its cosine coordinates are nonzero at $\tau = 0$).
2. **Finite dimensionality is a real loss.** That finitely many factors need not span the risks driving fixed-income markets is itself an established empirical theme [12]. The central argument of [13] — that hedging practice reflects maturity-specific risk which k -factor models misrepresent — is not answered by the completion. Whether the exchange of maturity-specific risk for a finite arbitrage-free completion is acceptable depends on whether one is pricing or forecasting, which is the dichotomy drawn in [13].
3. **The bond witness is structural.** Theorem 4.6 assumes frictionless continuous trading, strictly positive bond prices, and a generic finite panel, and it rebalances predictable cofactor exposures over a short interval. Transaction costs, position constraints, discrete observations, and market microstructure are not modeled.

Two questions stand open at these boundaries, and each is now posed exactly. The first is whether traded-bond processes are semimartingales under the fully infinite cylindrical dynamics. The repairing shift has infinite Cameron–Martin energy (Proposition 4.10), so the finite-dimensional repair does not extend to the infinite model as a change of measure. A bond-level or distribution-valued limit model — traded-bond semimartingales, a martingale measure, cross-maturity consistency — is where the infinite case would be decided. At the level of the factor laws the answer is known once the silent coordinates carry noise, the infinite laws being mutually singular (Proposition 4.11); it is the bond level that stays open, under both loadings. The second is whether the bilinear channel carries the short-end coefficient at positive vol-of-vol: a global conjugate-point problem, of which the uniform small-vol-of-vol region of [42] is the first piece.

What is settled is where the model’s finite sine truncations stand relative to the arbitrage-free class. They lie outside it at every resolution, by a margin that grows with resolution. The minimal repair is the doubled completion, whose cost is exact; and that cost is the HJM drift maturity-local volatility requires under $\mathbb{Q}$, its short-end value independent of the frozen intensity used to measure it.

References

- [1] Abi Jaber, E. and El Euch, O. (2019). Multifactor approximation of rough volatility models. *SIAM J. Financial Math.* **10**(2), 309–349.
- [2] Ahn, D.-H., Dittmar, R.F. and Gallant, A.R. (2002). Quadratic term structure models: theory and evidence. *Rev. Financial Studies* **15**(1), 243–288.

- [3] Albeverio, S., Marinelli, C. and Mastrogiacomo, E. (2020). Singular perturbations and asymptotic expansions for SPDEs with an application to term structure models. arXiv:2012.14510.
- [4] Andersen, L.B.G. and Piterbarg, V.V. (2007). Moment explosions in stochastic volatility models. *Finance Stoch.* **11**(1), 29–50.
- [5] Berman, A. and Plemmons, R.J. (1994). *Nonnegative Matrices in the Mathematical Sciences*. SIAM.
- [6] Björk, T. and Christensen, B.J. (1999). Interest rate dynamics and consistent forward rate curves. *Math. Finance* **9**(4), 323–348.
- [7] Björk, T. and Svensson, L. (2001). On the existence of finite-dimensional realizations for nonlinear forward rate models. *Math. Finance* **11**(2), 205–243.
- [8] Carmona, R. and Tehranchi, M. (2006). *Interest Rate Models: An Infinite Dimensional Stochastic Analysis Perspective*. Springer.
- [9] Christensen, J.H.E., Diebold, F.X. and Rudebusch, G.D. (2011). The affine arbitrage-free class of Nelson–Siegel term structure models. *J. Econometrics* **164**(1), 4–20.
- [10] Cheng, P. and Scaillet, O. (2007). Linear-quadratic jump-diffusion modeling. *Math. Finance* **17**(4), 575–598.
- [11] Cialenco, I. (2018). Statistical inference for SPDEs: an overview. *Stat. Inference Stoch. Process.* **21**(2), 309–329.
- [12] Collin-Dufresne, P. and Goldstein, R.S. (2002). Do bonds span the fixed income markets? Theory and evidence for unspanned stochastic volatility. *Journal of Finance* **57**(4), 1685–1730.
- [13] Cont, R. (2005). Modeling term structure dynamics: an infinite dimensional approach. *Int. J. Theor. Appl. Finance* **8**(3), 357–380. Preprint arXiv:cond-mat/9902018 (1999).
- [14] Dai, Q. and Singleton, K.J. (2000). Specification analysis of affine term structure models. *J. Finance* **55**(5), 1943–1978.
- [15] Da Prato, G. and Zabczyk, J. (2014). *Stochastic Equations in Infinite Dimensions*, 2nd ed. Cambridge University Press.
- [16] Dybvig, P., Ingersoll, J. and Ross, S. (1996). Long forward and zero-coupon rates can never fall. *J. Business* **69**, 1–25.
- [17] Duffee, G.R. (2002). Term premia and interest rate forecasts in affine models. *J. Finance* **57**(1), 405–443.
- [18] Duffie, D. and Kan, R. (1996). A yield-factor model of interest rates. *Math. Finance* **6**(4), 379–406.
- [19] Filipović, D. (2001). *Consistency Problems for Heath–Jarrow–Morton Interest Rate Models*. Springer LNM 1760.

- [20] Filipović, D. (2002). Separable term structures and the maximal degree problem. *Math. Finance* **12**(4), 341–349.
- [21] Filipović, D. and Teichmann, J. (2004). On the geometry of the term structure of interest rates. *Proc. R. Soc. A* **460**, 129–167.
- [22] Goldstein, R. (2000). The term structure of interest rates as a random field. *Rev. Financial Studies* **13**(2), 365–384.
- [23] Hansen, L.P. and Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *J. Political Economy* **99**(2), 225–262.
- [24] Hansen, L.P. and Jagannathan, R. (1997). Assessing specification errors in stochastic discount factor models. *J. Finance* **52**(2), 557–590.
- [25] Heath, D., Jarrow, R. and Morton, A. (1992). Bond pricing and the term structure of interest rates. *Econometrica* **60**, 77–105.
- [26] Hu, G.X., Pan, J. and Wang, J. (2013). Noise as information for illiquidity. *J. Finance* **68**(6), 2341–2382.
- [27] Hubalek, F., Klein, I. and Teichmann, J. (2002). A general proof of the Dybvig–Ingersoll–Ross theorem: long forward rates can never fall. *Math. Finance* **12**(4), 447–451.
- [28] Huebner, M. and Rozovskii, B.L. (1995). On asymptotic properties of maximum likelihood estimators for parabolic stochastic PDE’s. *Probab. Theory Related Fields* **103**(2), 143–163.
- [29] Kabanov, Yu.M. and Kramkov, D.O. (1998). Asymptotic arbitrage in large financial markets. *Finance Stoch.* **2**(2), 143–172.
- [30] Kakutani, S. (1948). On equivalence of infinite product measures. *Ann. of Math.* **49**(1), 214–224.
- [31] Karatzas, I. and Shreve, S.E. (1991). *Brownian Motion and Stochastic Calculus*, 2nd ed. Springer.
- [32] Karbach, S. (2026). Heat modulated affine stochastic volatility models for forward curve dynamics. arXiv:2409.13070.
- [33] Kennedy, D.P. (1994). The term structure of interest rates as a Gaussian random field. *Math. Finance* **4**(3), 247–258.
- [34] Leippold, M. and Wu, L. (2002). Asset pricing under the quadratic class. *J. Financial Quant. Anal.* **37**(2), 271–295.
- [35] Lyashenko, A. and Goncharov, Y. (2021–). Factor HJM framework. SSRN 3995533.
- [36] Musiela, M. (1993). Stochastic PDEs and term structure models. Preprint.
- [37] Nelson, C.R. and Siegel, A.F. (1987). Parsimonious modeling of yield curves. *J. Business* **60**(4), 473–489.

- [38] Rokhlin, D.B. (2008). Asymptotic arbitrage and numéraire portfolios in large financial markets. *Finance Stoch.* **12**(2), 173–194.
- [39] Santa-Clara, P. and Sornette, D. (2001). The dynamics of the forward interest rate curve with stochastic string shocks. *Rev. Financial Studies* **14**(1), 149–185.
- [40] Swanson, J. (2007). Variations of the solution to a stochastic heat equation. *Ann. Probab.* **35**(6), 2122–2159.
- [41] Walsh, J.B. (1986). An introduction to stochastic partial differential equations. In *École d'Été de Probabilités de Saint-Flour XIV*, Springer LNM 1180.
- [42] Cao, Z. and Huang, S. (2026). A uniform bounded-real criterion for quadratic–Gaussian term-structure existence. Companion manuscript.