Compatibility of Forecasts, Pricing and Physical Equilibrium of Interest Rates

Humped Affine Forward Volatility and Long-Forward Restrictions

Sachin Oza
Independent Researcher

18 November 2026

Abstract

Can a one-state representation of the forward-rate curve satisfy forecast consistency, arbitrage-free pricing, physical stationary equilibrium and long-forward restrictions simultaneously? We give a conditional compatibility classification for an exactly affine instantaneous-forward curve with a positive square-root short-rate diffusion. An interior hump in instantaneous-forward volatility selects a non-restoring affine risk-neutral drift but remains compatible with a broad family of restoring physical stationary laws, including non-CIR dynamics; neither pricing nor stationarity identifies a unique physical density or market price of risk.

A distinct and explicitly optional assumption—that the stationary physical expected forward curve is globally nondecreasing—implies that its physical mean short rate does not exceed the invariant mean of the long-forward measure, $\theta_P \le \theta_L$. On the humped branch, writing $u = \theta_P/\rho$ for the ratio to the asymptotic forward rate and $R > 1$ for the instantaneous-forward volatility hump ratio, the restriction is $u \le [1 - \sqrt{1 - 1/R}]/2$; equivalently, $0 < u < 1/2$ and $R \le [4u(1-u)]^{-1}$. A stationary-density-score interpretation clarifies how physical and pricing compatibility coefficients differ. Limited fitted-Treasury-yield diagnostics do not identify the pricing branch: a flat loading performs almost as well as the Riccati loading, while short-/long-tenor proxies cannot reject the mean inequality. The framework identifies which joint restrictions are demonstrable, which require additional geometry, and which physical features remain unidentified.

Keywords: affine forward curve; physical stationarity; instantaneous-forward volatility hump; square-root diffusion; market price of risk; long-forward measure; HJM compatibility.
JEL: C02, E43, G12.

1 Introduction

The central question is whether four descriptions of a single term-structure model—forecast representation, risk-neutral pricing, stationary physical equilibrium, and long-forward asymptotics—can be made mutually consistent. Their coexistence restricts the admissible class without necessarily identifying every part of the physical law; we distinguish deductions from optional geometry selections throughout.

The term structure of interest rates describes how interest rates vary with maturity. At any date the market supplies a forward-rate curve, but that curve is not fixed: the rates at different maturities move as economic conditions, expectations and risk compensation change. A natural question is whether the evolution of the entire curve can be described consistently by the stochastic behaviour of a single short interest rate.

The question becomes especially relevant if the short rate possesses a stationary probability distribution under the physical measure. By *stationary* we mean that the distribution can remain unchanged over time, even while individual realised rates continue to fluctuate. What, in that

setting, should be meant by an *equilibrium forward-rate curve*? And how can this interpretation coexist with the distinct requirements of arbitrage-free bond pricing and the observed geometry of forward-rate volatility?

Expected forward curves and the expected short rate

Consider a one-state model in which a given short rate determines a complete instantaneous-forward curve. Different possible short rates at a future date then generate different possible future curves. There are two natural ways to construct an expected future curve. One can evaluate the curve for each possible future short rate and average the resulting forward rates maturity by maturity. Or one can first take the expected future short rate and evaluate the forward curve at that expected value.

We adopt the forecasting selection that these two procedures agree. In notation introduced below, it is

$$\mathbb{E}_r^P[x(\tau, r_h)] = x\Big(\tau, \mathbb{E}_r^P[r_h]\Big)\,. \tag{1.1}$$

The condition says that a forecast of the short rate is sufficient to represent the expected *curve*, not merely its short-maturity boundary. Under appropriate smoothness and nondegeneracy assumptions, requiring this agreement across initial states, maturities and sufficiently short forecast horizons selects a curve that is *affine in the short-rate state*:

$$x(\tau, r) = c(\tau) + d(\tau)r. \tag{1.2}$$

Affinity is therefore motivated by an expected-value principle, rather than adopted solely to make bond pricing tractable. It does *not* mean that the forward curve is a straight line across maturities: both coefficients can vary nonlinearly with maturity.

The equilibrium curve implied by stationarity

Suppose that under the physical probability measure the short rate has an invariant distribution π_P and finite stationary mean $\theta_P = \mathbb{E}_{\pi_P}[r]$. Affinity then immediately gives

$$\mathbb{E}_{\pi_P}[x(\tau, r)] = x(\tau, \theta_P). \tag{1.3}$$

Thus the forward curve evaluated at the stationary mean short rate is exactly the average of the forward curves generated by the stationary distribution. This supplies a natural probabilistic interpretation of the *physical equilibrium forward curve*. It does not require realised curves to stop moving, a Gamma stationary distribution, or a particular physical drift function.

There is a distinction worth preserving. Physical stationarity does not, by itself, imply that a process *started* at its stationary mean has that same conditional expected short rate at every future horizon. The stronger requirement of constant-speed conditional-mean relaxation selects an affine physical drift under standard regularity assumptions; it is not a condition adopted here. We retain the expected-value consistency of the curve without imposing that stronger mean-relaxation condition. This allows the present work to examine nonlinear physical drifts without giving up the equilibrium-curve interpretation.

Pricing and the volatility problem

An equilibrium interpretation based on physical expectations is not yet a pricing theory. Bond prices must also satisfy no-arbitrage restrictions. In a one-state representation, Heath–Jarrow–Morton (HJM) consistency links the forward curve to the short-rate dynamics under the risk-neutral

probability measure. Under an exactly affine curve representation and the selected positive-rate boundary, these restrictions yield square-root short-rate variance.

We then ask for a further feature: an interior maximum in the *instantaneous-forward volatility loading*. In other words, the instantaneous sensitivity of a forward rate to the driving diffusion initially increases with maturity and subsequently decreases. This should not be confused with a hump in the forward-rate curve itself, the yield curve, or bond-return volatility.

The conventional *restoring risk-neutral* CIR branch cannot produce such a volatility hump: its maturity volatility loading declines monotonically. But this is not a theorem ruling out *physical* mean reversion, or even physical CIR dynamics. The pricing and physical measures need not assign the same drift to the short rate. We show that a non-restoring risk-neutral pricing drift produces a unique volatility hump while a broad class of physically stationary, restoring dynamics remains admissible. The market price of risk reconciles the two drifts; it need not be specified as a particular function in advance.

Finally, we investigate one further equilibrium-curve selection: that the *stationary physical expected forward curve* does not decline with maturity. This is not a consequence of stationarity and we do not conceal its strength. It gives an exact restriction on the physical stationary mean, with an interpretation through the drift of the long-forward measure. The restriction already exists on the zero-slope pricing branch and is strengthened, rather than created, by the volatility hump. Expressed in terms of the volatility hump and the long forward rate, the restriction has quantitative, potentially falsifiable implications without requiring reconstruction of the physical drift or market price of risk.

The four compatibility perspectives

The four perspectives play distinct roles. Forecast representation asks when averaging state-dependent curves commutes with averaging the underlying state. Risk-neutral pricing imposes the bond-pricing equation and fixes the maturity evolution of the affine pricing loadings. Physical stationarity specifies the possible invariant laws and associated physical drifts, without requiring stationarity under the ordinary risk-neutral measure. Long-forward asymptotics give an interpretation of the additional geometry imposed on the stationary expected curve. None of the four perspectives should be silently substituted for another, especially when moving between physical and pricing measures.

Selections, deductions and scope

The analysis therefore proceeds in a deliberate order. We first select expected-value consistency, which motivates state affinity. We separately select physical stationarity, which gives meaning to the average equilibrium curve. No-arbitrage pricing and the positive zero-variance floor then restrict the admissible diffusion. The volatility-hump requirement selects a branch of the risk-neutral pricing dynamics; physical stationarity determines a family of compatible physical drifts and market prices of risk. Only after these steps do we consider the *optional* nondecreasing-equilibrium-curve selection and derive its quantitative consequences.

Our aim is not a new pricing equation, a proof that observed rates are one-state diffusions, or an economic general-equilibrium derivation of the pricing kernel. The mathematical ingredients have extensive precedents in affine term-structure pricing (Cox et al., 1985; Duffie and Kan, 1996; Dai and Singleton, 2000), flexible market prices of risk (Duffee, 2002; Cheridito et al., 2007), stationary diffusion theory, and long-horizon pricing measures (Hansen and Scheinkman, 2009; Qin and Linetsky, 2016a, 2018). The contribution is to make the joint forecasting, equilibrium, volatility and pricing compatibility problem explicit, distinguish a restoring *physical* drift from a restoring *risk-neutral* drift, and show exactly what the additional equilibrium geometry commits us to.

Empirical calibration and monetary-policy intervention are kept separate from the mathematical results.

Empirical scope. The theory specifies instantaneous-forward volatility through $B'(\tau)$, whereas fitted yield curves must be differentiated to recover instantaneous forwards, often creating instability at long maturities. The integrated yield loading $B(\tau)/\tau$ avoids that differentiation. Our short historical illustration explicitly compares its fit against a flat-loading benchmark and uses imperfect short-/long-maturity proxies to convey the strength of the optional equilibrium selection. Neither calculation identifies the risk-neutral law, the physical stationary mean or the asymptotic forward rate.

The paper demonstrates both the strengths and the limitations of compatibility: it identifies restrictions that must hold under joint selections, as well as aspects of physical equilibrium that those selections cannot determine.

2 The state-to-curve map and the pricing equation

2.1 Notation and forecasting selection

For a time-homogeneous one-state diffusion let $r_t > 0$ be the short rate and $\tau \geq 0$ time to maturity. We write $P(\tau, r)$ for the time-zero value of a unit zero-coupon payment at horizon τ, with $P(0, r) = 1$, and define

$$x(\tau, r) = -\partial_\tau \log P(\tau, r), \qquad X(\tau, r) = \int_0^\tau x(s, r)\, \mathrm{d}s = -\log P(\tau, r). \tag{2.1}$$

As motivated in the introduction, "physical equilibrium" means an invariant probability law, not a general-equilibrium derivation of preferences or the pricing kernel. The short-end consistency condition is $x(0, r) = r$. To implement the expected-value forecasting selection, we use the state-affine instantaneous-forward field,

$$x(\tau, r) = c(\tau) + d(\tau)r, \qquad X(\tau, r) = A(\tau) + B(\tau)r, \tag{2.2}$$

where $A' = c$, $B' = d$, $A(0) = B(0) = 0$ and $B'(0) = 1$. No assumption is made that x is linear in maturity. We reserve $q(\tau, r) = \sigma(r)x_r(\tau, r)$ for instantaneous-forward volatility, $\ell(\tau) = B(\tau)/\tau$ for the zero-coupon yield loading, and $F(t, r)$ for the Feynman–Kac function in Appendix B. The coefficients $\mathcal{K}_X, \mathcal{K}_x$ introduced in Section 4.3 are distinct from any physical or pricing drift slopes κ_P, κ_Q.

For any physical conditional law with the requisite finite moment,

$$\mathbb{E}_r^P[x(\tau, r_h)] = x\big(\tau, \mathbb{E}_r^P[r_h]\big). \tag{2.3}$$

This formalises the expectation-commutation principle explained in the introduction. To see why this selection motivates affinity rather than merely accommodates it, temporarily allow a twice continuously differentiable function $x(\tau, r)$ and let $\sigma_P(r) > 0$ in the interior. For a fixed maturity and initial state, the physical generator and a Taylor expansion give, as $h \downarrow 0$,

$$\mathbb{E}_r^P[x(\tau, r_h)] - x\big(\tau, \mathbb{E}_r^P[r_h]\big) = \frac{1}{2}\sigma_P^2(r)\, \partial_{rr}x(\tau, r)\, h + o(h). \tag{2.4}$$

Indeed, the drift terms cancel between the two expansions. If expectation commutation holds for every interior state, maturity and sufficiently small h, the positive diffusion variance forces $\partial_{rr}x(\tau, r) = 0$, hence affinity on the connected state domain. This implication requires the regularity and moment conditions needed to interchange the generator and expectation; it is

not asserted for degenerate or irregular processes. Once affinity is selected, (2.3) is an algebraic identity and places no restriction on the physical drift. In contrast, the optional condition $\mathbb{E}_r^P[r_h] = \theta_P + (r-\theta_P)e^{-\kappa_P h}$ for all states and horizons implies $b_P(r) = \kappa_P(\theta_P - r)$ under standard generator regularity. We *do not* impose that additional condition in the main classification.

2.2 Affine pricing rigidity

Let under the ordinary risk-neutral measure

$$\mathrm{d}r_t = b_Q(r_t)\,\mathrm{d}t + \sigma(r_t)\,\mathrm{d}W_t^Q, \qquad v(r) = \sigma^2(r) > 0 \quad (r > 0). \tag{2.5}$$

The Feynman–Kac bond-pricing equation, or equivalently the one-state representation of HJM no-arbitrage, is

$$\partial_\tau P = b_Q(r)\partial_r P + \tfrac{1}{2}v(r)\partial_{rr}P - rP. \tag{2.6}$$

Inserting $P = e^{-A-Br}$ yields the identity

$$A'(\tau) + B'(\tau)r = r + B(\tau)b_Q(r) - \tfrac{1}{2}B(\tau)^2 v(r). \tag{2.7}$$

Proposition 2.1 (Affine state representation restricts pricing coefficients)**.** *Suppose* (2.7) *holds for all interior states and for two positive maturities at which B takes distinct, nonzero values. Then $v(r)$ and $b_Q(r)$ are affine functions of r. If v extends continuously to zero, $v(0) = 0$, and $v(r) > 0$ on $(0,\infty)$, then $v(r) = \beta r$ for some $\beta > 0$.*

Proof. At a maturity τ_i with $B_i := B(\tau_i) \neq 0$, write $A_i' := A'(\tau_i)$ and $B_i' := B'(\tau_i)$. Equation (2.7), divided by B_i, gives

$$b_Q(r) - \frac{1}{2}B_i v(r) = \frac{A_i'}{B_i} + \frac{B_i' - 1}{B_i}r.$$

The right-hand side is affine in r. Subtracting this identity at two maturities with $B_1 \neq B_2$ shows that $(B_2 - B_1)v(r)/2$ is affine, hence $v(r)$ is affine. Inserting that conclusion into either identity proves that $b_Q(r)$ is affine. The floor condition removes the constant term of v, and interior positivity forces a positive linear coefficient. □

The square-root diffusion is thus a consequence of exact one-state affine pricing and the stipulated *zero-variance* positive-rate boundary. Positivity alone would not select square-root diffusion, and affine models on different state domains can retain a nonzero variance intercept. The proposition also does not derive the sign of the risk-neutral drift slope.

3 Selecting an interior instantaneous-forward volatility hump

Start with the general affine risk-neutral drift $b_Q(r) = m + \alpha r$, where $m > 0$, $\alpha \in \mathbb{R}$, and $\sigma(r) = \sqrt{\beta r}$ with $\beta > 0$. The Riccati equation is $B' = 1 + \alpha B - \beta B^2/2$, and differentiation yields $B'' = (\alpha - \beta B)B'$. Since $B \geq 0$ and $B' > 0$ at every finite maturity, $\alpha \leq 0$ gives $B'' \leq 0$, with strict inequality away from the short boundary. If $\alpha > 0$, B crosses α/β exactly once below its positive limiting root and B'' changes sign from positive to negative. Thus a *strict interior hump selects* $\alpha > 0$ within this affine, positive-rate pricing class.

We henceforth write $\alpha = a > 0$, giving the selected branch

$$\boxed{b_Q(r) = m + ar, \qquad \sigma(r) = \sqrt{\beta r}}, \qquad m > 0,\ a > 0,\ \beta > 0. \tag{3.1}$$

For strict nonattainment of zero by the CIR-type diffusion started in the positive interior we assume $2m \geq \beta$; the weaker boundary conventions are not used in this paper. The Q diffusion

is nonexplosive and admits finite-horizon bond prices even though its undiscounted expectation grows.

Matching constant and linear coefficients in (2.7) gives

$$A' = mB, \qquad B' = 1 + aB - \tfrac{1}{2}\beta B^2, \qquad A(0) = B(0) = 0. \tag{3.2}$$

Therefore

$$x(\tau, r) = mB(\tau) + B'(\tau)r, \qquad q(\tau, r) = \sqrt{\beta r}\, B'(\tau), \tag{3.3}$$

where q denotes the *signed instantaneous diffusion coefficient of the forward rate at fixed remaining maturity*, not yield volatility, level variance or bond-return volatility. It is positive under our conventions.

Define

$$\gamma = \sqrt{a^2 + 2\beta}, \quad B_\infty = \frac{a+\gamma}{\beta}, \quad B_- = \frac{a-\gamma}{\beta} < 0. \tag{3.4}$$

The Riccati solution increases from zero to B_∞ and satisfies $B'(\tau) > 0$ at finite horizons. Its curvature is

$$B''(\tau) = (a - \beta B(\tau))B'(\tau). \tag{3.5}$$

Proposition 3.1 (A unique forward-volatility hump)**.** *The instantaneous-forward volatility $q(\tau, r)$, at each fixed state $r > 0$, rises from its short-maturity value, reaches exactly one interior maximum at the unique τ_* satisfying $B(\tau_*) = a/\beta$, and then decreases to zero. Its peak-to-short-boundary loading ratio is $1 + a^2/(2\beta)$.*

Proof. The bracket in (3.5) is positive for $B < a/\beta$ and negative for $B > a/\beta$. Since $0 < a/\beta < B_\infty$ and B is increasing, there is exactly one crossing. At the crossing, (3.2) gives $B' = 1 + a^2/(2\beta)$, while $B'(0) = 1$. □

If instead $b_Q(r) = m - \kappa_Q r$ with $\kappa_Q > 0$, then $B'' = -(\kappa_Q + \beta B)B' < 0$. Thus the no-hump conclusion for the restoring risk-neutral CIR branch and Proposition 3.1 are consistent. In particular, *physical* mean reversion does not enter the sign calculation.

3.1 Finite-horizon and long-horizon pricing are different questions

The ordinary Q conditional mean is

$$\mathbb{E}_r^Q[r_h] = \left(r + \frac{m}{a}\right) e^{ah} - \frac{m}{a}, \tag{3.6}$$

which diverges as $h \to \infty$. This does not invalidate bond pricing, because the latter concerns $\mathbb{E}^Q[e^{-\int_0^\tau r_s\,\mathrm{d}s}]$, not an undiscounted long-horizon expectation. In fact $B(\tau) \to B_\infty$, $A(\tau)/\tau \to mB_\infty$, and the long zero-coupon yield is finite and strictly positive, $\rho = mB_\infty$. As $\tau \to \infty$, the bond *price* tends to zero; it is the long yield and the state sensitivity that have finite limits. No stationary Q distribution is assumed or required.

4 Physical stationarity without selecting a physical drift

Under the physical probability measure let

$$\mathrm{d}r_t = b_P(r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^P, \tag{4.1}$$

where the function b_P is not prescribed. Fix the change-of-measure convention

$$b_Q(r) = b_P(r) + \sqrt{\beta r}\lambda(r), \tag{4.2}$$

so that $\mathrm{d}W^Q = \mathrm{d}W^P - \lambda(r_t)\,\mathrm{d}t$ under the corresponding local Girsanov relation.

Assumption 4.1 (Regular conservative physical equilibrium)**.** We work with a regular time-homogeneous diffusion on $(0,\infty)$ with volatility $\sqrt{\beta r}$, inaccessible endpoints, and a finite-mass invariant probability law with a strictly positive C^1 density π_P and finite first moment. Coefficients are locally regular enough for weak uniqueness and the scale–speed construction. Stationary-generator identities are invoked only with the corresponding moment, domain and vanishing-boundary-term conditions. The stationary current is zero *as a consequence* of the conservative, inaccessible-boundary, finite-speed-measure setting; it is not an additional economic closure.

For this conservative one-dimensional recurrent diffusion, the stationary Fokker–Planck current is constant and vanishes at the inaccessible boundaries. Thus the zero-current identity is

$$b_P(r)\pi_P(r) = \tfrac{1}{2}\frac{\mathrm{d}}{\mathrm{d}r}[\beta r \pi_P(r)], \tag{4.3}$$

which reconstructs the physical drift:

$$\boxed{b_P(r) = \frac{\beta}{2}\left[1 + r\partial_r \log \pi_P(r)\right].} \tag{4.4}$$

By (4.2), the associated market price of risk is then

$$\boxed{\lambda(r) = \frac{m + ar - \frac{\beta}{2}[1 + r\partial_r \log \pi_P(r)]}{\sqrt{\beta r}}.} \tag{4.5}$$

The statements do not mean that the density is observable without error, or that a historical sample spanning several policy regimes represents one time-homogeneous invariant law.

4.1 Scale, speed and boundary admissibility

For a positive density π normalised to integrate to one, define the candidate drift b_π by (4.4). Its scale density and speed density, up to positive multiplicative constants, are

$$s'_\pi(r) = \frac{1}{\beta r \pi(r)}, \qquad m'_\pi(r) = 2\pi(r). \tag{4.6}$$

Let $c \in (0,\infty)$ be arbitrary. The conditions

$$\int_0^c \frac{\mathrm{d}r}{r\pi(r)} = \infty, \qquad \int_c^\infty \frac{\mathrm{d}r}{r\pi(r)} = \infty \tag{4.7}$$

place both endpoints at infinite scale distance. Within the regular diffusion class on the open interval $(0,\infty)$, these conditions make both endpoints inaccessible to the minimal diffusion: a continuous path cannot hit an endpoint at which the scale function diverges. With finite speed mass, the resulting conservative diffusion is positive recurrent and has invariant density π. These are classical one-dimensional scale–speed results; see the sections on scale functions, speed measures, boundary classification and ergodicity in Rogers and Williams (2000, Chapter V). No reflecting or entrance-from-outside boundary extension is used. The finite first moment is required for the equilibrium-curve constructions later in the paper, not for positive recurrence itself.

Theorem 4.2 (Conditional stationary-density classification)**.** *Fix $m, a, \beta > 0$ as in* (3.1) *and assume $2m \geq \beta$. The specified Q diffusion is then strictly positive and conservative when started at an interior state. Let π be a strictly positive C^1 probability density on $(0,\infty)$ with*

$$\int_0^\infty \pi(r)\,\mathrm{d}r = 1, \qquad \int_0^\infty r\pi(r)\,\mathrm{d}r < \infty,$$

and satisfying both scale-divergence conditions (4.7)*. Define*

$$b_\pi(r) = \frac{\beta}{2}\left[1 + r\partial_r \log \pi(r)\right].$$

Assume that the SDE $\mathrm{d}r_t = b_\pi(r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^P$ *admits a weak solution, unique in law* up to the first exit *from* $(0,\infty)$*. No global conservativeness of this candidate P diffusion is assumed.*

Then the exit time is almost surely infinite, the P diffusion is strictly positive and positive recurrent, and its invariant probability measure is $\pi(r)\,\mathrm{d}r$*. Its scale and speed densities can be normalised as in* (4.6)*; the pointwise drift-change kernel relative to the specified Q diffusion is* (4.5)*. If P and Q start from the same interior state, or from mutually absolutely continuous initial probability measures on* $(0,\infty)$*, their laws on canonical continuous-path space are mutually absolutely continuous on* $\mathcal{F}_T$ *for every finite* T*.*

Conversely, within Assumption 4.1, any such regular stationary P diffusion has zero stationary current; hence its invariant density, drift and drift-change kernel necessarily obey (4.4)–(4.5)*. Positive recurrence implies that its scale function diverges at both endpoints, so* (4.7) *is necessary in the stated class.*

Proof. By the *definition* of b_π, the identity

$$b_\pi(r)\pi(r) = \frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}r}\left[\beta r \pi(r)\right]$$

holds algebraically; it does not presuppose that stationarity has been proved. Direct substitution into the scale and speed formulae gives $s'_\pi = 1/(\beta r \pi)$ and $m'_\pi = 2\pi$. Therefore the scale function diverges at both endpoints by (4.7), while the total speed mass is $\int_0^\infty 2\pi(r)\,\mathrm{d}r = 2$. For the minimal regular diffusion whose local existence and uniqueness were stipulated, the standard scale–speed recurrence and boundary criteria now establish *no exit in finite time*, positive recurrence and invariance of the normalised speed measure $\pi(r)\,\mathrm{d}r$ (Rogers and Williams, 2000, Chapter V). In particular, positivity and infinite lifetime are conclusions, not hypotheses about the candidate P law.

On the Q side, $2m \geq \beta$ prevents hitting zero and linear drift growth excludes explosion. On each compact $[1/n, n]$, the two drifts and the strictly positive diffusion coefficient give a bounded Girsanov kernel; its stopped exponential is therefore a true martingale. Weak uniqueness identifies the stopped transformed law, and Lemma B.1 in Appendix B passes from stopped equivalence to mutual equivalence on each canonical $\mathcal{F}_T$, because neither law exits $(0,\infty)$ by finite time. Equivalent initial distributions are covered by conditioning on the common initial-state variable. This does not imply equivalence on the infinite-horizon sigma field.

Conversely, for a regular stationary diffusion in Assumption 4.1, the stationary current

$$J(r) = b_P(r)\pi_P(r) - \tfrac{1}{2}\frac{\mathrm{d}}{\mathrm{d}r}\left[\beta r \pi_P(r)\right]$$

is constant and vanishes under the conservative recurrent no-boundary-flow setup. Thus $J \equiv 0$ recovers (4.4), and the Q-minus-P drift identity recovers (4.5). An invariant probability in this irreducible regular class gives positive recurrence and hence recurrence. The recurrence criterion forces the scale function to diverge at both endpoints, giving (4.7). □

Remark 4.3 (What the theorem does not identify). The pricing parameters and the existence of a volatility hump do not pick a unique π, b_P or λ. They describe a family. The density reconstruction determines the drift only for a chosen π and a fixed diffusion coefficient, with zero current and specified boundary behaviour. It is a familiar one-dimensional diffusion identity, not a new statistical identification device.

4.2 An equivalent risk-premium parameterisation

Substituting $b_P(r) = m + ar - \sqrt{\beta r}\lambda(r)$ into the stationary speed formula gives

$$\pi_P(r) = \frac{1}{Z} r^{2m/\beta - 1} \exp\left\{ \frac{2a}{\beta} r - \frac{2}{\sqrt{\beta}} \int_{r_*}^{r} \frac{\lambda(u)}{\sqrt{u}} \, \mathrm{d}u \right\}, \tag{4.8}$$

provided the candidate normalising constant Z is finite and the resulting density obeys the boundary and moment conditions. Formula (4.8) describes the same admissible class from the risk-premium side. It does not license arbitrary functions λ: the generated P law must be nonexplosive, stationary and finite-horizon equivalent to Q.

If $\lambda(r)$ stays bounded near zero, then $\pi_P(r) \sim C r^{2m/\beta - 1}$ as $r \downarrow 0$ (subject to local regularity). Thus $2m \geq \beta$ also makes the physical zero boundary inaccessible under this sufficient floor-premium selection. An unbounded premium is not automatically inadmissible; its boundary behaviour requires separate checking.

At the opposite tail, a globally bounded λ makes the exponential in (4.8) behave as $\exp[(2a/\beta)r + O(\sqrt{r})]$, preventing normalisation. Therefore a bounded risk premium at *all* large states is impossible in this stationary humped branch. A sufficient, but not necessary, tail condition is

$$\lambda(r) \geq \left(\frac{a}{\sqrt{\beta}} + \epsilon \right) \sqrt{r} \quad \text{for all sufficiently large } r, \qquad \epsilon > 0, \tag{4.9}$$

which gives a physically restoring upper drift bound $b_P(r) \leq m - \epsilon\sqrt{\beta}\, r$. Precise necessity is expressed by the integrability and scale conditions, not by the particular sufficient asymptotic inequality.

4.3 Compatibility coefficients and the stationary-density score

We write $\mathcal{K}_X$ for the integrated-price coefficient and $\mathcal{K}_x$ for the instantaneous-forward coefficient, reserving κ_P and κ_Q for drift slopes. The stationary-density reconstruction permits a physical-statistical interpretation of two coefficients that arise when the pricing compatibility equation is expressed in terms of the integrated-forward field and the instantaneous-forward field. These coefficients are *defined* here for comparison; neither is an additional constant or a separate equilibrium assumption.

Write $X(\tau, r) = -\log P(\tau, r)$, $x = X_\tau$, and define the integrated and instantaneous forward-volatility loadings by

$$G(\tau, r) = \sigma(r) X_r(\tau, r), \qquad q(\tau, r) = \sigma(r) x_r(\tau, r), \qquad \sigma(r) = \sqrt{\beta r}.$$

For $r > 0$ and $\tau > 0$, where X_r, G, x_r and q do not vanish, define state derivatives at fixed remaining maturity by

$$\mathcal{K}_X(\tau, r) = \frac{G_r}{X_r} - \frac{x - r + \frac{1}{2}G^2 - \lambda G}{G}, \tag{4.10}$$

$$\mathcal{K}_x(\tau, r) = \frac{\sigma q_r}{q} + \lambda(r). \tag{4.11}$$

Here the distinct symbols q (instantaneous-forward volatility) and G (bond-price volatility) are used throughout. The ratio G_r/X_r represents $\mathrm{d}G/\,\mathrm{d}X$ along the state coordinate at fixed maturity, and $\sigma q_r/q$ similarly represents $\mathrm{d}q/\,\mathrm{d}x$.

Proposition 4.4 (Stationary-density score and pricing-drift decomposition)**.** *Under the affine bond-pricing representation and the regular zero-current physical stationarity assumptions of Section 4, the quantities in* (4.10)–(4.11) *are independent of maturity and satisfy*

$$\boxed{\mathcal{K}_X(r) = -\frac{\sigma(r)}{2} \partial_r \log \pi_P(r),} \tag{4.12}$$

$$\boxed{\mathcal{K}_x(r) - \mathcal{K}_X(r) = \frac{b_Q(r)}{\sigma(r)}.} \tag{4.13}$$

In the humped square-root branch these read

$$\boxed{\mathcal{K}_X(r) = -\frac{\sqrt{\beta r}}{2}\partial_r \log \pi_P(r), \qquad \mathcal{K}_x(r) - \mathcal{K}_X(r) = \frac{m + ar}{\sqrt{\beta r}}.} \tag{4.14}$$

Proof. The pricing equation written under the physical measure, with the stated sign convention for λ, gives

$$x - r + \tfrac{1}{2}G^2 - \lambda G = b_P X_r + \tfrac{1}{2}\sigma^2 X_{rr}.$$

Since $G = \sigma X_r$, differentiation with respect to the state yields

$$\mathcal{K}_X = \sigma' - \frac{b_P}{\sigma} + \frac{\sigma}{2}\frac{X_{rr}}{X_r}.$$

Affine pricing has $X_{rr} = 0$, whereas the zero-current identity implies

$$b_P = \sigma\sigma' + \tfrac{1}{2}\sigma^2 \partial_r \log \pi_P.$$

This establishes (4.12). Similarly, $q = \sigma x_r$ and $x_{rr} = 0$ imply $\mathcal{K}_x = \sigma' + \lambda$. Consequently

$$\mathcal{K}_x - \mathcal{K}_X = \lambda + \frac{b_P}{\sigma} = \frac{b_Q}{\sigma},$$

which proves (4.13) and, on substitution of $\sigma = \sqrt{\beta r}$ and $b_Q = m + ar$, (4.14). □

Interpretation and limits. The function $\partial_r \log \pi_P$ is the *score* of the physical stationary density. Thus $\mathcal{K}_X$ is minus one half of that score weighted by the instantaneous short-rate diffusion coefficient. Where the stationary density increases, $\mathcal{K}_X < 0$; where it decreases, $\mathcal{K}_X > 0$; and at an interior stationary-density mode, $\mathcal{K}_X = 0$. It does not represent a coordinate-invariant restoring force or a constant physical mean-reversion speed. The difference $\mathcal{K}_x - \mathcal{K}_X$, by contrast, records the risk-neutral state drift normalised by the short-rate diffusion coefficient, with the physical drift and risk premium cancelling.

For a fixed diffusion coefficient, (4.12) can also be integrated:

$$\boxed{\pi_P(r) = C \exp\left[-2\int_{r_0}^{r} \frac{\mathcal{K}_X(u)}{\sigma(u)}\,\mathrm{d}u\right],} \tag{4.15}$$

where C is chosen to normalise the density, if possible, and the scale, speed and boundary conditions must still hold. This is a reparameterisation of the stationary-density reconstruction, not a new distributional identification theorem.

Under the optional locally bounded-premium selection $\lambda(r) = O(1)$ as $r \downarrow 0$, the affine identities additionally give

$$\boxed{\mathcal{K}_x(r) = \frac{\sqrt{\beta}}{2\sqrt{r}} + O(1), \qquad \mathcal{K}_X(r) = \frac{\beta/2 - m}{\sqrt{\beta r}} + O(1).} \tag{4.16}$$

Hence boundedness of the market price of risk does not entail boundedness of either compatibility coefficient. Equations (4.12)–(4.16) introduce no additional restriction on the admissible stationary family.

4.4 Stationary moments and their limits

Write $\theta_P = \mathbb{E}_{\pi_P}[r]$. Assuming r belongs to the appropriate extended generator domain with integrable drift,

$$\mathbb{E}_{\pi_P}[b_P(r)] = 0, \qquad \boxed{\mathbb{E}_{\pi_P}[\sqrt{\beta r}\lambda(r)] = m + a\theta_P > 0.} \tag{4.17}$$

The sign is a *weighted stationary mean* restriction and does not assert that $\lambda(r) > 0$ at every state. More generally, applying the stationary generator to r^n yields, if the requisite moments exist,

$$\mathbb{E}_{\pi_P}[r^{n-1}\sqrt{\beta r}\lambda(r)] = \left(m + \frac{n-1}{2}\beta\right)\mathbb{E}_{\pi_P}[r^{n-1}] + a\mathbb{E}_{\pi_P}[r^n], \quad n \geq 1. \tag{4.18}$$

These are standard stationary-generator identities. They restrict an admissible premium but do not determine its pointwise functional form.

Corollary 4.5 (Stationary instantaneous bond excess return)**.** *Under the integrability and zero-current stationarity assumptions used to derive* (4.17)*, the P-conditional instantaneous expected excess holding-return rate on a zero-coupon bond of remaining maturity τ is $B(\tau)\sigma(r)\lambda(r)$. Its stationary expectation is*

$$\boxed{\mathbb{E}_{\pi_P}[\textit{instantaneous excess bond-return rate at } \tau] = B(\tau)(m + a\theta_P).} \tag{4.19}$$

Proof. Under P, Itô's formula applied to $P(\tau, r) = e^{-A(\tau)-B(\tau)r}$ and the risk-neutral pricing PDE give a local expected bond holding-return rate $r + B(\tau)\sigma(r)\lambda(r)$, including the maturity roll-down. Taking the stationary expectation and applying (4.17) proves the result. □

The corollary concerns an instantaneous *expected holding-return premium* at a fixed maturity. It is neither the long-forward/physical-mean wedge $\rho - \theta_P$ nor, without additional assumptions, a conventional yield term premium or a finite-horizon expected excess return.

5 The physical equilibrium forward curve

The stationarity of r gives a time-independent physical mean curve,

$$x_e^P(\tau) := \mathbb{E}_{\pi_P}[x(\tau, r)] = mB(\tau) + \theta_P B'(\tau). \tag{5.1}$$

The stationary average curve equals the curve evaluated at the physical stationary mean, $x_e^P(\tau) = x(\tau, \theta_P)$; this is again the affine expected-value identity. It does *not* imply the conditional expectation of the curve remains unchanged when started from $r = \theta_P$ unless an additional physical conditional-mean condition holds.

We choose, as a separate equilibrium geometry selection, that x_e^P be nondecreasing over all remaining maturities. Differentiating (5.1) and using (3.5),

$$\partial_\tau x_e^P(\tau) = B'(\tau)\left[m + a\theta_P - \beta\theta_P B(\tau)\right]. \tag{5.2}$$

Since B' is strictly positive at finite maturities and $B \uparrow B_\infty$, the condition is exact.

Proposition 5.1 (Global equilibrium-curve geometry)**.** *For any physically stationary member of Theorem 4.2, the expected instantaneous-forward curve $x_e^P(\tau)$ is nondecreasing for all $\tau \geq 0$ if and only if*

$$\boxed{\theta_P \leq \frac{m}{\sqrt{a^2 + 2\beta}} = \frac{m}{\gamma}.} \tag{5.3}$$

If equality holds, its slope is positive at every finite maturity and tends to zero at infinity. This geometry condition is additional to physical stationarity.

Proof. Equation (5.2) is nonnegative for all maturities exactly when $m+a\theta_P \geq \beta\theta_P B_\infty = (a+\gamma)\theta_P$, which is equivalent to (5.3). □

The inequality places no pointwise sign requirement on b_P, and no equality between physical and pricing stationary distributions. It restricts only the first moment of the otherwise admissible physical density. The geometry selection is *additional* to the existence of the volatility hump: a positive a creates the hump, whereas the bound on θ_P enforces a globally nondecreasing expected curve. Its quantitative implications, including the part attributable to the hump and the part attributable to monotonicity alone, are derived in Section 7.

6 Long-horizon pricing and the long-forward measure

The divergence of $\mathbb{E}_r^Q[r_h]$ may seem unsettling despite finite bond prices. Pricing-eigenfunction methods clarify the distinction between ordinary Q expectations and long-horizon valuation probabilities (Hansen and Scheinkman, 2009; Qin and Linetsky, 2016a, 2018). Here the requisite eigenfunction is explicit.

Let $\mathcal{L}^Q f = (m + ar)f'(r) + \frac{1}{2}\beta r f''(r)$ and set

$$\phi(r) = e^{-B_\infty r}, \qquad \rho = mB_\infty. \tag{6.1}$$

The Riccati fixed-point identity $1 + aB_\infty - \frac{\beta}{2}B_\infty^2 = 0$ gives

$$(\mathcal{L}^Q - r)\phi = -\rho\phi. \tag{6.2}$$

Define the positive process

$$M_t = e^{\rho t - \int_0^t r_s \,\mathrm{d}s}\frac{\phi(r_t)}{\phi(r_0)}. \tag{6.3}$$

It is a local Q martingale with diffusion coefficient $-B_\infty\sqrt{\beta r_t}$. In this explicit affine model, its expectation is one at each finite horizon: the affine transform identity yields $\mathbb{E}_r^Q[e^{-\int_0^t r_s \,\mathrm{d}s}\phi(r_t)] = e^{-\rho t}\phi(r)$. Hence M_t is a true martingale and defines an equivalent measure $\mathbb{L}$ on finite horizons.

By the eigenfunction (Doob h-) transform, the L drift is

$$b_L(r) = b_Q(r) + \beta r\frac{\phi'(r)}{\phi(r)} = m + (a - \beta B_\infty)r = m - \gamma r. \tag{6.4}$$

This is a positive, restoring CIR diffusion with stationary Gamma density

$$\pi_L(r) \propto r^{2m/\beta-1}e^{-2\gamma r/\beta}, \qquad \boxed{\theta_L = \frac{m}{\gamma}.} \tag{6.5}$$

The recurrent L diffusion provides the long-forward measure for this explicit affine pricing model. This statement is an application of established factorisation theory, not a new probability measure.

The negative Riccati root $B_- < 0$ also produces a positive eigenfunction $\phi_-(r) = e^{-B_- r}$ and a pricing-operator eigenvalue mB_-. Its Doob transform has drift $m + \gamma r$ and is non-recurrent. The decreasing eigenfunction in (6.1), by contrast, selects the *recurrent* eigenfunction and its restoring transformed dynamics. Recurrent-eigenfunction uniqueness, rather than mere positivity of an eigenfunction, is the relevant long-forward selection; see Qin and Linetsky (2016a).

6.1 The finite-maturity forward-drift identity

For a fixed bond horizon τ, the τ-forward measure has time-dependent density obtained by taking the τ-bond as numeraire. At the initial date its short-rate drift is

$$b_{Q^\tau}(r;0) = b_Q(r) + \beta r \partial_r \log P(\tau, r) = m + ar - \beta r B(\tau). \tag{6.6}$$

At later calendar times $s < \tau$, replace $B(\tau)$ by $B(\tau - s)$. The initial-date qualification is essential: this is not a time-homogeneous finite-maturity forward-measure drift.

Proposition 6.1 (Physical equilibrium slope as a forward-measure drift)**.** *For every stationary physical law in Theorem 4.2 with mean θ_P,*

$$\boxed{\partial_\tau x_e^P(\tau) = B'(\tau)\, b_{Q^\tau}(\theta_P; 0).} \tag{6.7}$$

The forward-measure drift evaluated at θ_P is strictly decreasing as a function of bond horizon τ, and converges to $b_L(\theta_P) = m - \gamma\theta_P$. Hence

$$\boxed{\partial_\tau x_e^P(\tau) \geq 0 \text{ for all } \tau \geq 0 \quad \Longleftrightarrow \quad b_L(\theta_P) \geq 0.} \tag{6.8}$$

Proof. Differentiate $x_e^P = mB + \theta_P B'$ and use $B'' = (a - \beta B)B'$ to obtain $\partial_\tau x_e^P = B'[m + (a - \beta B)\theta_P]$, equal to (6.7) by (6.6). Since $\partial_\tau b_{Q^\tau}(\theta_P; 0) = -\beta\theta_P B'(\tau) < 0$ and $B(\tau) \uparrow B_\infty$, its limiting value is $m + (a - \beta B_\infty)\theta_P = m - \gamma\theta_P$. The positive factor B' and decreasing bracket prove the equivalence. □

Notice that the *same* Riccati state sensitivity $a - \beta B(\tau)$ controls the forward-volatility slope $B''(\tau)$ and the maturity-dependent component of the forward-measure drift. It changes sign precisely at the volatility maximum. The full drift $b_{Q^\tau}(\theta_P; 0)$ includes the positive intercept m, so its sign need not change there. This distinction prevents a misleading identification of the hump location with the equilibrium-curve turning point.

Corollary 6.2 (Physical versus long-forward equilibrium means)**.** *Within the classification of Theorem 4.2, the physical stationary expected instantaneous-forward curve is globally nondecreasing if and only if*

$$\boxed{\theta_P \leq \theta_L.} \tag{6.9}$$

No equality of P and L is required. Even equality of their means would not, in general, imply equality of their laws.

Proof. Proposition 6.1 gives the criterion $m - \gamma\theta_P \geq 0$. Since the recurrent long-forward law has mean $\theta_L = m/\gamma$, this is equivalent to the stated ordering. □

The changes of measure are related but distinct. Under our sign convention, the P-to-Q risk premium is $\lambda(r) = (b_Q - b_P)/\sqrt{\beta r}$. The corresponding P-to-L drift adjustment is

$$\lambda_L(r) = \frac{b_L(r) - b_P(r)}{\sqrt{\beta r}} = \lambda(r) - B_\infty\sqrt{\beta r}. \tag{6.10}$$

This is an identity, not an additional economic restriction. Setting $\lambda_L \equiv 0$ would force P=L; the present framework neither requires nor recommends it. The long-forward ordering is instead the interpretation of the already selected physical-equilibrium slope.

7 The observable content of the equilibrium selection

The comparison $\theta_P \le \theta_L$ is more than an alternative notation for the mean bound: Proposition 6.1 explains the *structural* role of the forward-measure drift. But the selection also imposes a strong and independently checkable constraint on the scale of the physical mean relative to the long forward rate.

7.1 Hump ratio, peak maturity and pricing parameters

Define the dimensionless, state-independent instantaneous-forward volatility ratio

$$R := \frac{q(\tau_*, r)}{q(0, r)} = 1 + \frac{a^2}{2\beta} > 1, \qquad \tau_* > 0. \tag{7.1}$$

Because $q(\tau, r) = \sqrt{\beta r}\, B'(\tau)$, the ratio does not depend on r and is not a ratio of yield volatilities or of realised finite-horizon forward changes. Appendix A gives the exact Riccati calculation

$$\gamma\tau_* = \log\left(\frac{\gamma + a}{\gamma - a}\right) = 2\,\mathrm{arcosh}\left(\sqrt{R}\right). \tag{7.2}$$

Write $s_R := \sqrt{1 - R^{-1}} \in (0, 1)$. The pair (R, τ_*) determines

$$\boxed{\gamma = \frac{2\,\mathrm{arcosh}(\sqrt{R})}{\tau_*}, \qquad a = \gamma s_R, \qquad \beta = \frac{\gamma^2}{2R}.} \tag{7.3}$$

Thus, under the exact time-homogeneous one-factor model and ideal observation of the *instantaneous* volatility loading, its peak height and location identify the pricing coefficients a and β; no physical drift estimation is used. The intercept m remains to be determined from the level of the curve. Define the long forward (and long zero-coupon yield) by

$$\rho := \lim_{\tau\to\infty} x(\tau, r) = mB_\infty = \frac{2m}{\gamma - a}. \tag{7.4}$$

Hence

$$\boxed{m = \frac{\rho}{2}(\gamma - a), \qquad \theta_L = \frac{m}{\gamma} = \frac{\rho}{2}(1 - s_R).} \tag{7.5}$$

These are exact *structural identification* statements, not claims that a smoothed market yield curve or a finite-window covariance estimator measures (R, τ_*, ρ) without uncertainty.

The restriction comes first from the optional *global monotonicity* selection, not from the volatility hump itself. More generally, for any real affine pricing slope a for which the same positive Riccati solution applies, $\gamma = \sqrt{a^2 + 2\beta}$ and $\rho = 2m/(\gamma - a)$ give

$$\boxed{\frac{\theta_P}{\rho} \le \frac{1}{2}\left(1 - \frac{a}{\gamma}\right).} \tag{7.6}$$

For $a < 0$ the bound is above one-half; for $a = 0$ it equals one-half; and for $a > 0$, precisely the humped-volatility branch, it is strictly below one-half. Thus the hump *tightens* a geometric restriction already present at $a = 0$, rather than creating the restriction. The ensuing formula in terms of R applies specifically to the $a > 0$ branch.

Corollary 7.1 (A falsifiable bound without risk-premium identification)**.** *If the physically stationary expected instantaneous-forward curve is globally nondecreasing and has a genuine interior instantaneous-forward volatility hump, then*

$$\boxed{\frac{\theta_P}{\rho} \le \frac{1}{2}\left(1 - \sqrt{1 - \frac{1}{R}}\right) < \frac{1}{2}.} \tag{7.7}$$

Consequently, $\rho > 2\theta_P$; for a specified ratio R, the bound is stronger still. In particular, the admissible stationary mean is less than approximately 29.6% of ρ when $R = 1.2$ and less than approximately 14.6% when $R = 2$.

Proof. Combine the equilibrium inequality $\theta_P \leq \theta_L$ with (7.5). Since $R > 1$, $s_R > 0$, giving strict inequality relative to 1/2. The numerical percentages follow by substitution. □

Corollary 7.2 (Inverse hump bound)**.** *Let $u = \theta_P/\rho > 0$. Under the same optional monotonicity selection, a genuine humped branch $R > 1$ requires*

$$\boxed{0 < u < \frac{1}{2}, \qquad 1 < R \leq \frac{1}{4u(1-u)}.} \tag{7.8}$$

Conversely, for $0 < u < 1/2$ this inequality is equivalent to (7.7). If $u = 1/2$, only the limiting $R = 1$ zero-slope branch can meet the bound; if $u > 1/2$, even that limiting branch is excluded. The inverse bound is not a bound on the non-humped $a < 0$ pricing branch.

Proof. For $0 < u < 1/2$, rearrange $2u \leq 1 - \sqrt{1 - 1/R}$ and square the nonnegative sides to obtain $1/R \geq 4u(1-u)$. If $u \geq 1/2$, the right-hand side of (7.7) is strictly less than 1/2 for $R > 1$. □

R	1.05	1.20	1.50	2.00
Maximum θ_P/ρ	39.1%	29.6%	21.1%	14.6%

7.2 Logical strength and potential rejection

A violation of (7.7), if its observables could be measured faithfully in the specified model, would reject the *joint* assumptions of the affine pricing branch, the volatility-hump measurement and the globally nondecreasing physical equilibrium curve. It would not separately disprove physical stationarity, the existence of the hump, or no-arbitrage. The bound is deliberately strong: as $R \downarrow 1$ it approaches 1/2 from below; as $R \to \infty$ it approaches zero. The one-half threshold at $a = 0$ originates in globally nondecreasing equilibrium geometry; the positive-slope pricing branch responsible for the hump tightens it. One must not soften it simply to improve an eventual empirical fit.

There is also a direct economic reading. Since $x_e^P(0) = \theta_P$ and $x_e^P(\infty) = \rho$, the humped branch requires $\rho - \theta_P > \theta_P$. Thus the long forward must stand more than one physical mean short rate above the stationary short-end average. This is a large *long-forward/physical-mean wedge*. It resembles a long-horizon term-compensation comparison, but it is not, without additional assumptions about risk adjustment and matching horizons, identical to a conventional term premium. Failure of this severe restriction in actual data is a serious possibility, and would challenge the joint selections rather than the mere mathematical coexistence of physical stationarity and humped pricing volatility. The stationary expected instantaneous excess bond-return rate is instead the precise quantity in Corollary 4.5, $B(\tau)(m + a\theta_P)$; its relation to a quoted yield term premium requires a separate horizon and expectations decomposition.

An empirical comparison would require a defensible estimate of the stationary physical short-rate mean and a long-forward-rate proxy, while distinguishing instantaneous forward volatility from fitted-maturity change volatility. Monetary-policy changes can invalidate the common stationary regime. These are substantial measurement challenges, not additional theoretical conditions. Crucially, the quantitative implication is derived *without* reconstructing b_P or λ.

8 Constructive compatibility examples

8.1 Physical CIR and a humped pricing branch

The simplest construction reveals why the original incompatibility is measure-specific. Let

$$m = a = \beta = 1, \qquad \mathrm{d}r_t = (1 - 3r_t)\,\mathrm{d}t + \sqrt{r_t}\,\mathrm{d}W_t^P, \tag{8.1}$$

with risk-neutral dynamics $\mathrm{d}r_t = (1 + r_t)\,\mathrm{d}t + \sqrt{r_t}\,\mathrm{d}W_t^Q$. Their market price of risk is $\lambda(r) = 4\sqrt{r}$. The physical mean is $\theta_P = 1/3$, while the long-forward mean is $\theta_L = 1/\sqrt{3}$. Hence $\theta_P < \theta_L$. The physical process is strictly positive, ergodic and ordinary CIR; the risk-neutral process is non-restoring. Nevertheless, $B' = 1 + B - \frac{1}{2}B^2$ produces a unique forward-volatility maximum at $B = 1$. The expected physical equilibrium curve rises over every finite maturity.

The example is in the established extended-affine risk-premium class. More generally, if $b_P(r) = m - \kappa_P r$ and $\lambda(r) = \lambda_1\sqrt{r}$, then $b_Q(r) = m + (\sqrt{\beta}\lambda_1 - \kappa_P)r$, so

$$\boxed{a = \sqrt{\beta}\lambda_1 - \kappa_P.} \tag{8.2}$$

Thus a positive pricing-drift slope is compatible with restoring physical CIR precisely when $\sqrt{\beta}\lambda_1 > \kappa_P$. The example is not a new class of pricing kernels; the pertinent risk-premium sign freedom is directly related to Duffee (2002) and Cheridito et al. (2007). Its role is to refute the inference that *physical* CIR dynamics alone block the hump. The no-hump restriction arises under *risk-neutral* restoration.

8.2 A non-Gamma stationary family with nonlinear physical drift

For a fixed pricing triple (m, a, β) choose $k \geq \gamma$ and $\delta \geq 0$, and consider

$$\pi_\delta(r) = \frac{1}{Z_\delta} r^{2m/\beta - 1} \exp\left[-\frac{2k}{\beta}r - \delta r^2\right]. \tag{8.3}$$

Equation (4.4) gives

$$b_{P,\delta}(r) = m - kr - \beta\delta r^2, \tag{8.4}$$

while (4.2) gives

$$\lambda_\delta(r) = \frac{a + k}{\sqrt{\beta}}\sqrt{r} + \sqrt{\beta}\,\delta r^{3/2}. \tag{8.5}$$

All these laws satisfy the positive-boundary and recurrence conditions when $2m \geq \beta$. For $\delta = 0$, the stationary mean is $m/k \leq m/\gamma$. The density for $\delta > 0$ is an exponential tilt of that baseline by the decreasing weight $e^{-\delta r^2}$. Chebyshev's association inequality, applied to the increasing function r and the decreasing function $e^{-\delta r^2}$, gives $\mathrm{Cov}(r, e^{-\delta r^2}) < 0$, and hence $\theta_{P,\delta} < m/k \leq m/\gamma$ for $\delta > 0$. Every member consequently meets the global equilibrium-curve selection. The $\delta > 0$ members have nonlinear physical drift and non-Gamma stationary densities, while preserving identical affine pricing and identical humped volatility maturity loading.

For $\delta > 0$, the $r^{3/2}$ component in (8.5) takes the risk premium outside the conventional extended-affine square-root specification used in the preceding CIR example. Its admissibility is instead ensured by the regular positive diffusion construction and the finite-horizon localisation argument in Lemma B.1; no global boundedness or Novikov assumption is made.

This family illustrates the exact scope of the classification: infinitely many distinct physical laws correspond to the same risk-neutral affine term-structure geometry. Without physical evidence or additional forecasting assumptions they cannot be differentiated by the pricing curve alone.

9 Forecasting, covariance and empirical meaning

9.1 What affinity guarantees, and what it does not

The state-affine map supplies (2.3) for every horizon and state whenever $\mathbb{E}_r^P[r_h]$ exists. By itself it does not prescribe how that mean evolves. Under the nonlinear family (8.4), the first-moment equation is

$$\frac{\mathrm{d}}{\mathrm{d}h}\mathbb{E}_r^P[r_h] = m - k\mathbb{E}_r^P[r_h] - \beta\delta\mathbb{E}_r^P[r_h^2], \tag{9.1}$$

which is not a closed constant-speed relaxation equation for $\delta > 0$. In particular, the stationary mean state need not be preserved in *conditional* expectation when starting from that state: $\mathbb{E}_{\theta_P}^P[r_h]$ need not equal θ_P at every h. A paper that imposes this stronger preservation must retain extra conditions and will generally select a narrower physical family.

Thus there are at least three distinct statements: (i) affine curve evaluation commutes with conditional expectation, (ii) the physical stationary average curve equals the mean-state curve, and (iii) starting exactly at that state gives a constant conditional mean curve at every future horizon. The first two hold in our framework; the third is not claimed outside appropriately selected physical dynamics.

9.2 Covariance restrictions from exact one-state affinity

At any fixed date and state, instantaneous forward innovations satisfy

$$\mathrm{d}x_t(\tau)\big|_{\text{diffusion}} = B'(\tau)\sqrt{\beta r_t}\,\mathrm{d}W_t^P, \tag{9.2}$$

with the same diffusion coefficient under Q. Hence the instantaneous conditional cross-tenor covariance is

$$\frac{\mathrm{d}\langle x(\tau_i), x(\tau_j)\rangle_t}{\mathrm{d}t} = \beta r_t B'(\tau_i)B'(\tau_j), \tag{9.3}$$

which has rank one if the loading vector is nonzero. Under physical stationarity, the instantaneous covariance averaged over the stationary law has scale $\beta\theta_P$ and the same maturity-product form. The stationary *level* covariance is likewise rank one: $\mathrm{Cov}_{\pi_P}(x(\tau_i, r), x(\tau_j, r)) = B'(\tau_i)B'(\tau_j)\,\mathrm{Var}_{\pi_P}(r)$, when the variance exists. Exact one-state affinity makes strong predictions; an empirical dominant principal component is not equivalent to exact rank one.

For nonlinear physical drift, the temporal covariance decay need not have one common exponential rate, even though contemporaneous cross-sectional covariance remains rank one. This is important when interpreting apparent changes in physical persistence: covariance geometry and time-lagged dynamics are different diagnostics. Under linear restoring P drift, conditional mean persistence is exponential; under general P drift, its form is not prescribed.

9.3 Integrated affine geometry: the zero-coupon yield as an observable

The previous covariance restriction concerns instantaneous forwards. For empirical use there is an important alternative: the price itself depends on an *integral* of the forward curve, rather than its maturity derivative. Define the continuously compounded zero-coupon yield at maturity $\tau > 0$ by

$$y(\tau, r) = -\frac{\log P(\tau, r)}{\tau} = \frac{A(\tau)}{\tau} + \frac{B(\tau)}{\tau}r, \qquad \ell(\tau) := \frac{B(\tau)}{\tau}. \tag{9.4}$$

The same affine pricing geometry therefore implies, at fixed maturity and fixed model parameters,

$$\boxed{y(\tau, r_{t+h}) - y(\tau, r_t) = \ell(\tau)(r_{t+h} - r_t),} \tag{9.5}$$

for any observation interval h, including when the physical drift is nonlinear. Equation (9.5) is a state-map identity, *not* an approximation based on instantaneous trading. Its conditional instantaneous diffusion loading is

$$q_y(\tau, r) = \sqrt{\beta r}\,\ell(\tau) = \frac{1}{\tau}\int_0^\tau q_x(s, r)\,\mathrm{d}s, \qquad q_x(\tau, r) = \sqrt{\beta r}\,B'(\tau). \tag{9.6}$$

The equality is a same-state, same-instant relation between *model diffusion loadings*; it is not an equality between sample standard deviations of yield changes and instantaneous forward-rate changes. The integrated loading avoids differentiating an estimated yield curve in maturity, which can magnify long-tenor fitting noise.

Proposition 9.1 (Integrated Riccati yield-loading restriction)**.** *In the time-homogeneous affine pricing model, let $\ell(\tau) = B(\tau)/\tau$. Then for every $\tau > 0$,*

$$\tau\ell'(\tau) + \ell(\tau) = B'(\tau), \qquad \tau\ell''(\tau) + 2\ell'(\tau) = \big(a - \beta\tau\ell(\tau)\big)\big(\tau\ell'(\tau) + \ell(\tau)\big), \tag{9.7}$$

with $\ell(0+) = 1$ and $\ell'(0+) = a/2$. For a fixed finite set of maturities, the covariance matrix of yield innovations is proportional to $\ell\ell^{\mathsf{T}}$ and hence has rank one. A dominant empirical principal component need not satisfy the nonlinear differential restriction in (9.7)*, even when it closely resembles some fitted ℓ.*

Proof. Differentiate $B(\tau) = \tau\ell(\tau)$ once and twice, then apply the Riccati identity $B'' = (a-\beta B)B'$. Taylor expansion of $B' = 1+aB-\beta B^2/2$ at zero gives $B(\tau) = \tau+(a/2)\tau^2+O(\tau^3)$, so $\ell'(0+) = a/2$. Finally $\mathrm{d}y(\tau, r_t)|_{\text{diffusion}} = \ell(\tau)\sqrt{\beta r_t}\,\mathrm{d}W_t$ at every tenor. □

This proposition separates three increasingly demanding empirical statements. First, a single maturity loading might explain most yield movements. Second, that loading might be approximated by the specific Riccati family $B(\tau)/\tau$. Third, the full time-homogeneous one-state model would require *exact* cross-tenor rank one, fixed loadings and the stated square-root state dependence, together with the remaining pricing and boundary conditions. Success at the first two levels does not establish the third. In particular, a loading fit alone may fail to discriminate $a > 0$ from $a \le 0$; it does not identify the physical drift, the pricing measure or the market price of risk.

9.4 Empirical scope: a diagnostic, not a pricing test

The integrated Riccati loading supplies a more stable observable than maturity-differentiated fitted instantaneous forwards, but good one-factor fit is not identification of the pricing branch. Appendix D compares the loading against both a flat maturity loading and an unrestricted first principal component using rolling Svensson-fitted U.S. Treasury curves. In the full sample the Riccati fit explains 91.19% of the fitted yield-change variation, only 0.63 percentage points more than the parameter-free flat benchmark; the positive and nonpositive pricing-drift branches are not distinguishable descriptively. The 64.9% full-sample short-/long-tenor proxy ratio is not a rejection statistic: under the selected monotonic mean curve, the analogous true stationary-curve ratio is an upper bound on θ_P/ρ, whereas a rejection needs a lower bound. The nearly indistinguishable flat and Riccati loading fits provide no persuasive identification of the humped branch. Neither comparison validates exact rank one, the diffusion law, or HJM pricing. The paper's compatibility results do not depend on these diagnostics.

9.5 Monetary-policy intervention and partial identification

The invariant density is a property of the complete physical stochastic system. For a policy-sensitive short rate it may reflect markets, monetary-policy reaction functions, administered-rate

changes and regime transitions. A decomposition $b_P = b_{\text{market}} + b_{\text{policy}}$ is not identified by observing a single rate path. Policy can also move volatility, change market prices of risk and add effective states, so a time-homogeneous one-state diffusion may not be empirically adequate across regimes.

Consequently this paper does not infer a market-only physical drift from a historical sample, equate changes in observed forward-rate persistence with changes in a specific κ_P, or regard covariance loading stability as proof of stable risk-neutral pricing parameters. The fitted-yield illustration above is descriptive rather than an empirical identification or calibration of the full model. The theoretical results remain conditional statements about the assumed autonomous one-state model. Any eventual decisive exercise would need independently grounded bond prices or yield observations, an observed short-rate state, tests of exact cross-maturity restrictions, and a treatment of regime variation before reconstructing λ from an estimated density.

9.6 No resurrection of a physical conservation-of-return pricing rule

A physical rolling-maturity transport or conservation-of-return hypothesis cannot substitute for the fixed-calendar-maturity HJM pricing equation. The compatibility identities here are derived from the bond-pricing PDE, which already encodes the relevant no-arbitrage restrictions, and a distinct physical stationary generator. Any proposed physical conservation law would require an explicit coordinate and change-of-measure argument before it could be compared with the pricing drift. No such extra law is imposed in this paper.

10 Relation to prior literature and scope of contribution

The affine rigidity argument belongs to the affine diffusion and yield-factor literature (Cox et al., 1985; Duffie and Kan, 1996; Dai and Singleton, 2000); the presence of a Riccati fixed point and its curvature consequences is algebraically straightforward. The separation of physical and risk-neutral dynamics and broader choices of the market price of risk have extensive precedents (Duffee, 2002; Cheridito et al., 2007). Nonparametric interest-rate drift and density approaches are also established (Aït-Sahalia, 1996; Stanton, 1997). Scale–speed construction, recurrence and boundary admissibility use standard one-dimensional diffusion theory (Rogers and Williams, 2000, Chapter V). The eigenfunction and recurrent long-forward measure follow established long-term-factorisation methods (Hansen and Scheinkman, 2009; Qin and Linetsky, 2016a, 2018, 2016b).

The humped *instantaneous-forward volatility* studied here is not the same concept as a humped *yield curve*; shape restrictions for one-factor affine yields have an independent literature (Keller-Ressel and Steiner, 2008). Nor is it a claim that HJM models generally require square-root diffusion: the restriction follows from our one-state, state-affine, zero-variance-floor selections. General HJM volatility surfaces need not have the one-state affine form.

The paper's specific contribution is to make the strengths and limitations of compatibility explicit across four representations of the same model: forecast averaging, arbitrage-free pricing, physical stationarity, and long-forward structure. Their simultaneous use selects a particular pricing geometry and restricts admissible physical equilibria, but does not uniquely identify the physical drift or invariant density. The stationary-density-score interpretation of the compatibility coefficient makes the remaining freedom transparent; it is an interpretation of established identities, not an additional selection principle. This is an organisational and conditional mathematical contribution, not a claim to a new elementary stochastic differential equation or priority over the constituent literatures. In particular, $\theta_P \leq \theta_L$ holds under the precise additional equilibrium-curve geometry selection used here, not as a universal physical-pricing inequality.

11 Conclusion

A one-state term structure can be read simultaneously as a forecast representation, a risk-neutral pricing system, a stationary physical stochastic law, and a long-forward asymptotic structure. Asking that these descriptions coexist restricts the model sharply; it does not by itself determine a unique physical equilibrium or certify the model empirically.

Exact affine instantaneous-forward rates provide a valuable expected-value closure but, with an inaccessible positive floor and vanishing floor variance, severely restrict the pricing diffusion. The resulting square-root variance admits a unique hump in instantaneous-forward volatility when the affine risk-neutral drift has a positive state slope. This branch is inconsistent with a stationary *ordinary* Q short-rate law, but not with finite-horizon arbitrage-free bond pricing or physical stationarity.

Under the regular stationary physical equilibrium class, the stationary density reconstructs a physical drift and the associated market price of risk. In the affine branch, the master-equation compatibility coefficient is minus one half the volatility-weighted score of that density, while its difference from the corresponding instantaneous-forward coefficient is the risk-neutral short-rate drift divided by volatility; these are interpretations, not additional restrictions. Scale, speed, moments and finite-horizon equivalence delimit the admissible family. The physical drift need not be CIR; conversely, physical CIR itself can coexist with the hump. A globally nondecreasing stationary expected forward curve is governed by the drift of the finite-maturity forward measure evaluated at the physical stationary mean. Its limiting sign is the long-forward drift $b_L(\theta_P) = m - \gamma\theta_P$, giving the mean ordering $\theta_P \le \theta_L$. The observable restriction $\theta_P/\rho \le \frac{1}{2}(1 - \sqrt{1 - 1/R}) < 1/2$, equivalently $R \le [4u(1-u)]^{-1}$ for $u = \theta_P/\rho \in (0, 1/2)$, reveals how demanding this particular global geometry selection is: global monotonicity already imposes a one-half bound when $a = 0$, and a genuine hump tightens it. Economically, the humped branch demands a long-forward/physical-mean wedge greater than the physical mean itself; this is not identified unconditionally with a standard term premium. The stationary instantaneous bond-return premium $B(\tau)(m + a\theta_P)$ provides a separate expected-return interpretation, not an identification of the yield term premium. Neither equality of P and L nor any prescribed physical drift is needed.

These are conditional mathematical compatibility results, not an empirical identification of a one-state interest-rate law. The integrated Riccati yield loading $B(\tau)/\tau$ is observable in principle, but the illustrative fitted-yield comparison in Appendix D offers no compelling identification of the positive-a pricing branch: a flat loading fits almost as well. Short-/long-maturity proxies likewise do not identify the physical stationary mean or long forward rate. Policy regimes, additional yield factors and measurement effects remain separate empirical questions. The illustrative 64.9% proxy ratio cannot reject the mean restriction, while the negligible improvement over a flat loading weighs against treating the hump as empirically identified. Thus the permissive compatibility theorem and the restrictive optional monotonicity selection must be evaluated on their own logical terms. The hump is a useful, sharply discriminating application of the framework; the more general lesson is which conclusions follow from jointly consistent representations, which depend on optional economic or geometric selections, and which remain unidentified.

A Additional Riccati and pricing details

Factoring the Riccati polynomial gives $B' = \frac{\beta}{2}(B_\infty - B)(B - B_-)$, where $B_- < 0 < B_\infty$. Separation of variables yields

$$\frac{B(\tau) - B_-}{B_\infty - B(\tau)} = \frac{-B_-}{B_\infty} e^{\gamma\tau}. \tag{A.1}$$

Equivalently, with $K = (a+\gamma)/(\gamma-a) > 1$,

$$B(\tau) = B_\infty \frac{e^{\gamma\tau}-1}{e^{\gamma\tau}+K}. \tag{A.2}$$

The formula implies $B(0)=0$, $B'(0)=1$, $B \uparrow B_\infty$ and $B' \downarrow 0$ after its unique maximum. The peak maturity can be expressed without numerical optimisation: solving $B(\tau_*) = a/\beta$ gives

$$\tau_* = \frac{1}{\gamma}\log\left(\frac{\gamma+a}{\gamma-a}\right). \tag{A.3}$$

Since $R = 1 + a^2/(2\beta) = \gamma^2/(\gamma^2-a^2)$, we have $a/\gamma = \sqrt{1-1/R}$ and $\gamma\tau_* = 2\,\mathrm{arctanh}(a/\gamma) = 2\,\mathrm{arcosh}(\sqrt{R})$, as used in Section 7. The explicit peak expression agrees with the implicit characterisation $B(\tau_*) = a/\beta$. The bond pricing solution $P = e^{-A-Br}$, $A = m\int_0^\tau B(s)\,\mathrm{d}s$, is strictly between zero and one for positive rates and horizons.

B Finite-horizon measure equivalence and eigenfunction martingale

Lemma B.1 (From local Girsanov equivalence to finite-horizon equivalence)**.** *Let P and Q be the laws on canonical continuous-path space $C([0,\infty),(0,\infty))$, equipped with the canonical filtration $(\mathcal{F}_t)$, of two weakly unique one-dimensional diffusions started at the same $r_0 > 0$. Suppose both have the same continuous, strictly positive diffusion coefficient in the interior, both drifts are locally regular, and neither diffusion reaches 0 or explodes to infinity in finite time. For all sufficiently large integers n, let*

$$\tau_n = \inf\{t \geq 0 : r_t \notin (1/n, n)\}.$$

Then for every finite T, the restrictions of P and Q to $\mathcal{F}_T$ are mutually absolutely continuous.

Proof. Fix $T < \infty$ and take n large enough that $r_0 \in (1/n, n)$. Under P put $\lambda(r) = [b_Q(r) - b_P(r)]/\sigma(r)$ and define

$$Z_t^{(n)} = \exp\left\{\int_0^{t\wedge\tau_n} \lambda(r_s)\,\mathrm{d}W_s^P - \frac{1}{2}\int_0^{t\wedge\tau_n}\lambda(r_s)^2\,\mathrm{d}s\right\}, \quad 0 \leq t \leq T. \tag{B.1}$$

The kernel is bounded on $[1/n, n]$; consequently $Z^{(n)}$ is a true P-martingale on $[0,T]$. The probability defined by $Z_T^{(n)}\,\mathrm{d}P$ makes $W^Q_{t\wedge\tau_n} = W^P_{t\wedge\tau_n} - \int_0^{t\wedge\tau_n}\lambda(r_s)\,\mathrm{d}s$ Brownian up to stopping, and changes the stopped drift to b_Q. Weak uniqueness identifies its stopped coordinate law with the Q-law stopped at τ_n. Since the stopped density is strictly positive, the stopped laws are mutually absolutely continuous on $\mathcal{F}_{T\wedge\tau_n}$. This argument can equally be run in reverse with kernel $-\lambda$ under Q.

Put $E_n = \{\tau_n > T\}$. For $A \in \mathcal{F}_T$, $A \cap E_n$ is determined by the stopped path and belongs to $\mathcal{F}_{T\wedge\tau_n}$. If $Q(A) = 0$, stopped equivalence gives $P(A\cap E_n) = 0$. Both candidate laws are assumed conservative with continuous paths in $(0,\infty)$, so $P(E_n) \to 1$ and $Q(E_n) \to 1$. It follows that $P(A) = 0$. Interchanging P, Q proves the converse and hence $P|_{\mathcal{F}_T} \sim Q|_{\mathcal{F}_T}$.

The laws here start at a common interior state. Mutually absolutely continuous initial probability measures on $(0,\infty)$ also suffice after conditioning on the initial state, but a deterministic initial state and a nondegenerate stationary initial distribution do *not* define equivalent path laws at time zero. No equivalence on the infinite-horizon path sigma-field is asserted. In particular, a global Novikov condition is not required: the proof uses true stopped martingales and the separate conservativeness of both candidate diffusions. □

The lemma is a standard localisation argument rather than a new measure-change theorem. In the present application, the candidate P law is constructed by the scale-speed criterion and the Q law is the nonexplosive affine square-root diffusion. Their inaccessibility assumptions and local regularity verify the lemma's hypotheses.

To justify the eigenfunction martingale without assuming that every formal PDE solution is an expectation, note that $\phi(r) = e^{-B_\infty r}$ is bounded between zero and one for $r > 0$, and the killing factor $e^{-\int_0^t r_s \,\mathrm{d}s}$ is also bounded by one. The Feynman–Kac expectation

$$F(t,r) = \mathbb{E}_r^Q\left[e^{-\int_0^t r_s \,\mathrm{d}s}\phi(r_t)\right]$$

is therefore well defined and bounded. The bounded-solution Feynman–Kac uniqueness principle for the conservative square-root diffusion identifies it with the explicit bounded classical solution $F(t,r) = e^{-\rho t}\phi(r)$, obtained from the constant Riccati root B_∞ and intercept rate $\rho = mB_\infty$. Consequently $\mathbb{E}_r^Q[M_t] = 1$ for every finite t. Equivalently, a localisation construction of the h-transform gives the conservative restoring diffusion $\mathrm{d}r_t = (m - \gamma r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^L$; the absence of boundary exit prevents loss of mass when stopping is removed. Both routes establish that the positive local martingale M is a true finite-horizon martingale. Neither proves $P = L$ or infinite-horizon equivalence.

C Proof audit: assumptions, consequences and nonclaims

1. **Selected:** time homogeneity, one continuous short-rate state, exact state-affine forwards and the positive zero-variance floor. **Derived:** state-affine pricing drift and variance, and square-root variance on the chosen state domain.
2. **Selected:** a strict interior instantaneous-forward volatility hump. **Derived:** the sign of the affine Q drift slope and the hump location on the resulting Riccati branch.
3. **Selected:** regular conservative physical stationarity with inaccessible boundaries and finite speed mass. **Derived:** zero stationary current, the density–drift–risk-premium correspondence and weighted moment restrictions.
4. **Selected:** globally nondecreasing physical stationary expected forwards. **Derived:** forward-measure drift sign criterion, long-forward mean ordering, and the restrictive observable inequality linking θ_P/ρ to hump ratio R.
5. **Not selected:** equality P=L, bounded ordinary Q mean at infinite horizons, proportional physical conditional-mean relaxation, a policy-adjusted short rate, a particular market-price-of-risk family or a physical conservation-of-return pricing axiom.

This audit is intended to make the logical status of each assertion visible. The theorem is a classification inside the stated boundary and regularity class, not a characterization of every possible HJM model.

D Limited diagnostics from fitted Treasury yields

Purpose. This appendix provides an in-sample descriptive diagnostic, not a test of the risk-neutral pricing martingale or the physical stationary law. To illustrate the identification problem without treating it as verification of the theory, we compared $\ell(\tau)$ with historical U.S. Treasury *fitted yields*. The input is a monthly rolling-curve workbook derived from the Federal Reserve nominal Treasury curve estimates (`https://www.federalreserve.gov/data/nominal-yield-curve.htm`): 538

month-end observations from September 1981 to June 2026, each averaging the preceding 60 daily Svensson-fitted zero-coupon curves. These are model-fitted yields, not independent bond-price observations or instantaneous conditional volatilities. Within each sample, three-month changes across maturities 1–15 years are demeaned, and the fractions of pooled squared changes captured by one-factor projections are compared. An unrestricted first principal component gives the best unrestricted one-factor fit; the Riccati model restricts the loading to $B(\tau)/\tau$; the flat comparator fixes the loading at $\ell(\tau) \equiv 1$.

Table 1: Explained pooled variation in fitted zero-coupon yield changes, 1–15 years (overlapping three-month changes; percent).

Sample	Flat $\ell \equiv 1$	$a > 0$ Riccati	$a \leq 0$ Riccati	PC1
1981–1999	92.77	93.67	93.70	93.72
2000–2009	82.59	85.90	85.81	86.03
2010–2019	88.11	93.85	88.10	94.34
2020–2026	91.72	92.43	92.51	92.55
Full sample	90.56	91.19	91.19	91.24

Boundary-fit note. The flat loading $\ell(\tau) \equiv 1$ is the limit of $B(\tau)/\tau$ as $a \to 0$ and $\beta \to 0$. Because the fitted Riccati family imposes $\beta > 0$, this limiting shape need not be attained by a finite admissible parameter pair. The 88.10% constrained $a \leq 0$ figure versus 88.11% for the flat benchmark in 2010–2019 is a negligible numerical/boundary-fit difference, not an economic finding.

In the full sample the positive-a Riccati loading improves the parameter-free flat benchmark by only 0.63 percentage points and cannot be distinguished descriptively from the $a \leq 0$ branch; the fitted positive-a optimum approaches zero. Because the Riccati shape has two fitted parameters whereas the flat shape has none, this 0.63-percentage-point in-sample increment is the observed gain from the additional flexibility, not independent evidence for the pricing model. Thus a 91% loading fit is not strong evidence for the pricing model. The 2010–2019 interval differs, but a zero-lower-bound policy regime, forward guidance and other policy measures may compress short-maturity yield changes and mimic the positive-a loading shape. This is an alternative explanation, not evidence identifying policy as the cause. Moreover, the empirical maturity grid begins at one year, so it never observes the model's short-end derivative $\ell'(0+) = a/2$.

We also provide a *proxy stress test* of the additional global monotonicity selection. The inequality on the humped branch constrains the *stationary physical mean* θ_P relative to the *asymptotic forward* ρ. The workbook measures neither. We therefore compare historical averages of the fitted one-year zero-coupon yield and fifteen-year instantaneous forward, which are not substitutes for those model objects:

Table 2: Illustrative short-/long-maturity proxy comparison; averages of fitted curves, not estimates of θ_P/ρ.

Sample	Mean 1Y yield (%)	Mean 15Y forward (%)	Ratio (%)
1981–1999	7.08	9.06	78.1
2000–2009	3.06	5.75	53.2
2010–2019	0.77	3.85	19.9
2020–2026	2.86	3.88	73.5
Full sample	4.16	6.41	64.9

Rounding note. Ratios use the unrounded sample means; the two rate columns are displayed to two decimal places. Thus, for example, the reported 2020–2026 ratio is 73.5%, whereas division of

the displayed 2.86% by 3.88% gives 73.7%.

The full-sample proxy ratio is 64.9%, visibly above one half, but this does *not* reject the theory. If the stationary mean curve is globally nondecreasing, its one-year average zero-coupon yield is at least θ_P and its fifteen-year instantaneous forward is at most ρ. For positive rates this makes the ratio of those two *population stationary-curve quantities* an *upper bound* on $u = \theta_P/\rho$; rejecting an upper bound on u requires a defensible *lower* bound on u. The reported cross-regime sample averages are not even guaranteed to estimate a common stationary curve. The more informative negative diagnostic is that the positive-a Riccati loading offers little full-sample improvement over the flat loading and does not distinguish the pricing-drift branches. The period 2010–2019 has the lowest proxy ratio while also yielding the greatest descriptive advantage for positive a; policy-induced differences in short-maturity variation are a plausible competing explanation.

A more direct, but still model-dependent, empirical exercise would fit the *stationary mean instantaneous-forward curve* $\bar{x}(\tau) = mB(\tau) + \theta_P B'(\tau)$ jointly over maturities, with the four parameters (m, a, β, θ_P), subject to the exact Riccati relation and positivity. Such a fit would estimate $u = \theta_P/(mB_\infty)$ and permit direct comparison with (7.8); the smooth theoretical mean forward curve avoids differentiating the *fitted theoretical function*, but constructing observed instantaneous forwards from fitted yields may still introduce differentiation and long-tenor noise. Identification, policy-regime stability and independent validation of the stationary mean would remain necessary.

Finally, the one-state theorem predicts *exact* rank-one yield changes with fixed pricing parameters. The roughly 9% residual variation in the pooled 1–15-year projection is therefore outside the exact model, not noise that the model itself explains. Overlapping 60-day averages and three-month differences, in-sample fitting, Svensson smoothing, and the use of raw rather than short-rate-volatility-normalised changes further limit inference. We make no pricing-validation or formal hypothesis-test claim. These comparisons are deliberately restricted illustrations; the compatibility theorem does not rely on them.

Declaration of generative AI assistance

During the development and preparation of this manuscript, the author used OpenAI ChatGPT for exploratory mathematical reasoning, derivation and algebra checks, assistance with computational diagnostics, and drafting, restructuring and editing the exposition. AI assistance extended beyond language correction and included suggestions for mathematical relationships and their presentation. The author critically reviewed and independently re-derived the principal mathematical results, evaluated and revised the AI-assisted material, and takes full responsibility for the mathematical claims, references, empirical qualifications and final manuscript. The use of AI does not constitute independent validation or peer review.

References

Rogers, L. C. G., and Williams, D. (2000). *Diffusions, Markov Processes and Martingales, Vol. 2: Itô Calculus*, 2nd ed., Chapter V (stochastic differential equations, diffusion theory, and one-dimensional boundary criteria). Cambridge University Press.

Aït-Sahalia, Y. (1996). Nonparametric pricing of interest rate derivative securities. *Econometrica*, 64(3), 527–560. `https://www.nber.org/papers/w5345`.

Cheridito, P., Filipović, D., and Kimmel, R. L. (2007). Market price of risk specifications for affine models: theory and evidence. *Journal of Financial Economics*, 83(1), 123–170. `https://doi.org/10.1016/j.jfineco.2005.09.008`.

Cox, J. C., Ingersoll, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2), 385–407.

Dai, Q., and Singleton, K. J. (2000). Specification analysis of affine term structure models. *Journal of Finance*, 55(5), 1943–1978.

Duffee, G. R. (2002). Term premia and interest rate forecasts in affine models. *Journal of Finance*, 57(1), 405–443. https://doi.org/10.1111/1540-6261.00426.

Duffie, D., and Kan, R. (1996). A yield-factor model of interest rates. *Mathematical Finance*, 6(4), 379–406. https://doi.org/10.1111/j.1467-9965.1996.tb00123.x.

Hansen, L. P., and Scheinkman, J. A. (2009). Long-term risk: an operator approach. *Econometrica*, 77(1), 177–234. https://doi.org/10.3982/ECTA6761.

Keller-Ressel, M., and Steiner, T. (2008). Yield curve shapes and the asymptotic short rate distribution in affine one-factor models. *Finance and Stochastics*, 12(2), 149–172. https://doi.org/10.1007/s00780-007-0059-z.

Qin, L., and Linetsky, V. (2016a). Positive eigenfunctions of Markovian pricing operators: Hansen–Scheinkman factorization, Ross recovery, and long-term pricing. *Operations Research*, 64(1), 99–117. https://doi.org/10.1287/opre.2015.1449.

Qin, L., and Linetsky, V. (2016b). Long-term factorization of affine pricing kernels. *Working paper*. https://arxiv.org/abs/1610.00778.

Qin, L., and Linetsky, V. (2018). Long-term factorization in Heath–Jarrow–Morton models. *Finance and Stochastics*, 22, 621–641. https://doi.org/10.1007/s00780-018-0365-7.

Stanton, R. (1997). A nonparametric model of term structure dynamics and the market price of interest rate risk. *Journal of Finance*, 52(5), 1973–2002. https://doi.org/10.1111/j.1540-6261.1997.tb02748.x.