A Compatibility Framework for Interest-Rate Term Structures

From Physical Forecasts and Stochastic Equilibrium to Arbitrage-Free Pricing

Sachin Oza

Independent Researcher

Consolidated manuscript — 18 November 2026

First assembled edition, for mathematical and editorial review

Prefatory note

This monograph develops, step by step, a conditional compatibility framework for one-state interest-rate term structures. Its purpose is to make the mathematics accessible to readers familiar with calculus, differential equations and probability without presuming specialist training in stochastic finance. It draws on the companion research paper, *Compatibility of Forecasts, Pricing and Physical Equilibrium: Humped Affine Forward Volatility and Long-Forward Restrictions* (Version 16, 18 November 2026). The chapters are assembled from the separately prepared teaching drafts. Their mathematical conclusions remain conditional on the hypotheses stated in the relevant sections; the empirical diagnostics are illustrative, not a validation of the full model.

Contents

Prefatory note **ii**

1 The Interest-Rate Term Structure as a Stochastic System **1**

1.1 Why a term structure is more than a collection of rates 1

1.2 A bond price as the primary observable 2

1.2.1 A promised payment at a future date 2

1.2.2 Zero-coupon yields 2

1.3 Instantaneous forward rates 2

1.3.1 Definition by maturity differentiation 2

1.3.2 Recovering forwards from yields 3

1.3.3 The short end 3

1.4 Two maturity coordinates, one evolving market 3

1.4.1 A simple example: movement without repricing 4

1.5 Realised curves, forecasts and prices are different objects 4

1.5.1 Why an observed forward is not automatically a forecast 5

1.6 Representing an entire curve using one state 5

1.6.1 The one-state hypothesis 5

1.6.2 State affinity is not maturity linearity 5

1.6.3 The integrated map 6

1.7 Two ways of forecasting a curve 6

1.7.1 Why affinity is sufficient 6

1.7.2 Why the converse requires hypotheses 7

1.8 Physical equilibrium is a probability law, not a motionless curve 7

1.8.1 Stationarity of the state distribution 7

1.8.2 A subtle but important non-implication 7

1.8.3 The equilibrium curve need not be increasing 8

1.9 The difference between physical and pricing descriptions 8

1.10 The hierarchy: selections, deductions and unresolved quantities 8

1.11 What Chapter 1 has established—and what remains free 9

2 Stochastic Dynamics, Itô Calculus and the Generator **11**

2.1 From a changing rate to a random process . . . 11
2.2 Brownian motion: three properties that matter . . . 12
2.2.1 Why ordinary calculus is not enough . . . 12
2.3 Deriving Itô's formula from a Taylor expansion . . . 12
2.3.1 Worked example: the square of the short rate . . . 13
2.3.2 Worked example: an exponential function . . . 13
2.4 The infinitesimal generator . . . 14
2.4.1 Why conditional expectations produce the generator . . . 14
2.5 From a single state to expected future curves . . . 14
2.5.1 The special simplification under state affinity . . . 15
2.6 An illustrative restoring diffusion: no CIR assumption required . . . 15
2.7 The generator under a change of probability measure . . . 16
2.8 What Chapter 2 has established—and what remains free . . . 16

3 Forecast Consistency and the Affine State-to-Curve Map 18
3.1 The problem: two ways to forecast one curve . . . 18
3.2 Why a nonlinear curve generally fails the test . . . 19
3.3 Why an affine map works for every physical forecast law . . . 19
3.4 The converse: when forecast closure forces affinity . . . 20
3.5 Examples showing why the hypotheses matter . . . 21
3.5.1 A deterministic process . . . 21
3.5.2 Equality at only one maturity . . . 21
3.5.3 Approximate rather than exact closure . . . 22
3.6 Affinity of the curve does not impose affine physical drift . . . 22
3.7 A stronger forecast condition would select mean-reverting affine drift . . . 22
3.8 Physical equilibrium as a separate averaging operation . . . 23
3.9 Forward curves, yields and bond prices under state affinity . . . 23
3.10 Physical forecasts are not risk-neutral prices . . . 24
3.11 What Chapter 3 has established—and what remains free . . . 24

4 Arbitrage-Free Bond Pricing and the Riccati Equations 26
4.1 What changes when we move from forecasting to pricing? . . . 26
4.2 The risk-neutral short-rate generator . . . 27
4.3 The exponential-affine representation: what is already known? . . . 27
4.4 Substitution into the pricing equation, without skipped steps . . . 28
4.5 Why two maturities constrain the pricing dynamics . . . 29
4.6 What selects square-root diffusion? . . . 29
4.7 A careful look at the short-maturity boundary . . . 30
4.8 The Riccati equation is not yet a volatility-hump selection . . . 30
4.9 Why rolling maturity must not be confused with a held bond . . . 31

4.10 A second route to the same pricing restrictions . . . 31
4.11 What Chapter 4 has established—and what remains free . . . 32

5 The Instantaneous-Forward Volatility Hump 34
5.1 The question inherited from Chapter 4 . . . 34
5.2 From bond prices to instantaneous-forward volatility . . . 35
5.3 Why the Riccati equation determines the hump . . . 35
5.3.1 The restoring risk-neutral branch . . . 36
5.3.2 The non-restoring risk-neutral branch . . . 36
5.4 How large is the hump? . . . 36
5.5 Finding the hump maturity step by step . . . 37
5.6 A transparent numerical illustration . . . 38
5.7 The risk-neutral drift is non-restoring—what does that mean? . . . 38
5.8 How physical mean reversion can coexist with the hump . . . 39
5.9 What the hump does—and does not—identify . . . 39
5.10 What Chapter 5 has established—and what remains free . . . 40

6 Physical and Risk-Neutral Measures 41
6.1 Why two probability measures are necessary . . . 41
6.2 A small-time picture of a change of measure . . . 41
6.3 Girsanov's theorem with an explicit sign convention . . . 42
6.4 Why the generator changes but the local variance does not . . . 43
6.5 Local forecast differences versus finite-horizon forecasts . . . 43
6.6 A worked example: restoring P and non-restoring Q . . . 44
6.7 Physical equilibrium and the mean premium identity . . . 45
6.8 Why the equilibrium forward curve initially rises . . . 45
6.9 What this chapter establishes—and what remains free . . . 46

7 Expected Excess Bond Returns and Physical Forecasts 47
7.1 Why bond returns are the missing connection . . . 47
7.2 The return on a bond held through calendar time . . . 48
7.3 The pointwise forecast–return compatibility identity . . . 49
7.3.1 Why curvature must cancel . . . 49
7.3.2 A cross-maturity restriction . . . 49
7.4 Linking bond returns to the forward-curve spread . . . 50
7.5 Averaging away from equilibrium . . . 50
7.6 The stationary equilibrium relation . . . 51
7.7 What would identify the physical stationary distribution? . . . 51
7.7.1 Physical CIR as a worked example, not a requirement . . . 52
7.7.2 Stationary expected-return distributions across maturities . . . 52

7.8 The empirical and logical limits . . . 52
7.9 What we have established—and what remains free . . . 52

8 Physical Equilibrium and Stationary Distributions 54
8.1 Why equilibrium is a probability distribution, not a constant rate . . . 54
8.2 From the generator to the forward Kolmogorov equation . . . 55
8.3 Probability current and conservation of probability . . . 55
8.4 Deriving the zero-current stationary density . . . 56
8.5 Formal solutions, recurrence and the zero boundary . . . 57
8.6 Worked example I: CIR and the Gamma density . . . 57
8.7 Worked example II: a non-CIR physical equilibrium . . . 58
8.8 Stationary moments and the meaning of equilibrium . . . 59
8.9 The stationary-density score and the compatibility coefficients . . . 59
8.10 What we have established—and what remains free . . . 60

9 The Physical Equilibrium Forward Curve 61
9.1 From a stationary short rate to an expected forward curve . . . 61
9.2 What equilibrium means at the short and long boundaries . . . 62
9.3 Deriving the maturity slope . . . 63
9.4 Why physical stationarity does not select monotonicity . . . 63
9.5 Risk compensation and the equilibrium forward spread . . . 64
9.6 A preview of the long-forward measure . . . 65
9.7 A quantitative consequence of the volatility hump . . . 66
9.8 Two physical CIR examples: same pricing hump, different equilibrium geometry . . . 67
9.9 What this chapter establishes, and what remains free . . . 67

10 Long-Horizon Pricing and the Long-Forward Measure 68
10.1 Three different questions, and three different measures . . . 68
10.2 The long-maturity Riccati fixed point . . . 69
10.3 Why a pricing eigenfunction appears . . . 69
10.4 From an eigenfunction to a change of probability measure . . . 70
10.5 The restoring short-rate dynamics under the long-forward measure . . . 71
10.6 Why the recurrent eigenfunction matters . . . 71
10.7 Finite-maturity forward measures and the limiting drift . . . 71
10.8 The equilibrium-curve slope and the long-forward mean . . . 72
10.9 The point of maximum forward volatility is not the point of maximum equilibrium forwards . . . 73
10.10 Changes of measure and what remains unidentified . . . 73
10.11 What this chapter establishes, and what remains free . . . 74

11 What Compatibility Determines, and What It Does Not 75

11.1 From coexistence to identification . . . 75
11.2 The hierarchy, revisited without compressing its assumptions . . . 76
11.3 What pricing geometry determines . . . 76
11.4 What pricing geometry cannot determine: two physical laws . . . 77
11.4.1 A restoring CIR physical law . . . 77
11.4.2 A nonlinear restoring physical law . . . 78
11.5 What additional information can identify the physical drift? . . . 78
11.5.1 From physical drift to stationary density . . . 79
11.5.2 Forecast evidence offers a second route . . . 79
11.6 Why equilibrium-curve information alone is not enough . . . 79
11.7 What the model makes observable and falsifiable . . . 80
11.7.1 Cross-tenor rank one . . . 80
11.7.2 The yield loading is more directly accessible . . . 80
11.7.3 The geometric mean bound . . . 80
11.8 A map of what is fixed, free and additionally observable . . . 81
11.9 What this chapter establishes, and what remains free . . . 81

12 Empirical Measurement and Tests of Compatibility 82
12.1 The purpose of empirical work . . . 82
12.2 Why yields are easier to observe than instantaneous forwards . . . 83
12.3 From a common factor to a Riccati maturity loading . . . 83
12.4 What the fitted U.S. Treasury diagnostics actually show . . . 84
12.5 Why the stationary-mean proxy cannot reject monotonicity . . . 85
12.6 Why the volatility hump is unusually hard to observe . . . 85
12.7 Expected excess returns offer another identification route . . . 86
12.8 A practical sequence of increasingly demanding tests . . . 87
12.9 Regime dependence, falsification and the limits of one state . . . 87
12.10What this chapter establishes, and what remains free . . . 88

13 Conclusions and Open Questions 89
13.1 Returning to the original question . . . 89
13.2 The hierarchy in one view . . . 89
13.3 The hump and the apparent conflict with equilibrium . . . 90
13.4 Risk compensation provides the compatibility mechanism . . . 91
13.5 Physical equilibrium is flexible, but not unconstrained . . . 91
13.6 The optional long-forward selection . . . 92
13.7 What the framework unifies—and what it does not . . . 93
13.8 Open questions and a disciplined next stage . . . 93
13.9 Closing perspective . . . 94

References **95**

Chapter 1

The Interest-Rate Term Structure as a Stochastic System

An interest-rate curve is observed at one date but describes payments at many future dates. This elementary fact creates several distinct mathematical questions: what does the curve measure; how does it change as calendar time advances; what does it tell us about the future short rate; and which restrictions follow from the absence of arbitrage? This chapter constructs the objects needed to state those questions precisely. It introduces zero-coupon bonds, instantaneous forwards, rolling versus fixed maturity, state-to-curve maps, physical forecasting and probabilistic equilibrium. It then states the compatibility programme without presupposing its conclusions. No stochastic calculus is required; the necessary machinery is developed in later chapters.

1.1 Why a term structure is more than a collection of rates

An interest rate is often presented as a single number. A bank account might pay a stated rate over the next month; a government bond might promise a payment in ten years. But a market at any one date contains prices for payments over many different horizons. The term structure is a way of describing those prices consistently.

Its mathematical interest lies in the combination of two dimensions. One is *calendar time*: the date on which we observe the market. The other is *maturity*: the date at which a future payment is due. As time advances, the market can reprice every maturity, even if no bond has yet matured. Moreover, a bond's remaining life declines simply because the calendar advances. These are separate effects and must not be confused.

Our central question is whether a single stochastic state, such as the instantaneous short rate, can describe a complete term structure in a manner consistent with physical forecasts, no-arbitrage prices, expected bond returns and stationary physical behaviour. The answer will be conditional: some properties will be selected as assumptions; others will follow mathematically; several economically interesting quantities will remain unidentified.

Guiding question. If the same short-rate process is used to describe the evolving curve, to forecast interest rates, to price bonds, and to define long-run equilibrium, what mathematical restrictions ensure that these descriptions do not contradict one another?

1.2 A bond price as the primary observable

1.2.1 A promised payment at a future date

At calendar date t, let

$$P(t, T) \tag{1.1}$$

be the price of a default-free zero-coupon bond that pays one unit of currency at the fixed future date $T > t$. Throughout this chapter we use idealised continuously compounded, default-free nominal claims; complications such as collateral conventions, liquidity, credit risk and market segmentation are outside the introductory model.

The price is dimensionless when the maturity payment is one unit of account. A value $P(t, T) = 0.90$, for example, means that a payment of one at T costs 0.90 at t. The bond price, rather than a quoted yield, is our primitive observable: yields and forward rates will be defined from it.

Define the *time to maturity*

$$\tau = T - t. \tag{1.2}$$

For a bond observed today with three years remaining, $\tau = 3$ years. Tomorrow its maturity date T will be unchanged, but its remaining time τ will be slightly shorter.

When no confusion arises, we write the same price as

$$\mathcal{P}_t(\tau) := P(t, t + \tau). \tag{1.3}$$

This is simply a change of coordinates; it does not define a new security.

1.2.2 Zero-coupon yields

For $\tau > 0$, define the continuously compounded zero-coupon yield $y_t(\tau)$ by

$$\mathcal{P}_t(\tau) = \exp\bigl[-\tau y_t(\tau)\bigr], \qquad y_t(\tau) = -\frac{\log \mathcal{P}_t(\tau)}{\tau}. \tag{1.4}$$

This is an identity defining a yield quote. No statistical or dynamic model is being imposed. Its interpretation is the constant continuously compounded rate that reproduces the observed discount factor over the interval $[t, t + \tau]$.

If a three-year discount factor is $e^{-0.12}$, its continuously compounded yield is $0.12/3 = 4\%$ per year. The price itself does not tell us that the short rate will be 4% at any particular future date.

1.3 Instantaneous forward rates

1.3.1 Definition by maturity differentiation

Assuming the discount curve is sufficiently smooth in maturity, define the instantaneous forward rate by

$$x_t(\tau) = -\partial_\tau \log \mathcal{P}_t(\tau). \tag{1.5}$$

We shall use $x_t(\tau)$ for the *realised curve at observation time* t. The adjective *instantaneous* refers to an infinitesimal future accrual interval around maturity $t + \tau$, not to an assertion that the future rate is known today.

Integrating equation (1.5) from zero to τ, and using $\mathcal{P}_t(0) = 1$, gives

$$\log \mathcal{P}_t(\tau) - \log \mathcal{P}_t(0) = -\int_0^\tau x_t(s)\,\mathrm{d}s, \tag{1.6}$$

$$\therefore \qquad \mathcal{P}_t(\tau) = \exp\left[-\int_0^\tau x_t(s)\,\mathrm{d}s\right]. \tag{1.7}$$

Combining (1.4) and (1.7) yields

$$\boxed{y_t(\tau) = \frac{1}{\tau}\int_0^\tau x_t(s)\,\mathrm{d}s.} \tag{1.8}$$

Thus a zero-coupon yield is the maturity average of instantaneous forwards between the observation date and the bond's maturity.

An elementary but consequential distinction. The forward curve is a derivative of the logarithm of the bond-price curve. The yield curve is an average of the forward curve. Neither is, by definition, the physical expected future short-rate curve.

1.3.2 Recovering forwards from yields

Equation (1.8) implies $\tau y_t(\tau) = \int_0^\tau x_t(s)\,\mathrm{d}s$. Differentiation gives

$$\boxed{x_t(\tau) = y_t(\tau) + \tau\partial_\tau y_t(\tau).} \tag{1.9}$$

This formula matters empirically. Observed bond yields are typically available at a finite number of maturities; constructing instantaneous forwards therefore involves fitting and differentiating a curve. Differentiation can amplify fitting noise and sensitivity to the chosen curve family. We will distinguish theoretical forward-rate statements from tests implemented using fitted yields.

1.3.3 The short end

The instantaneous short rate r_t is the idealised rate earned by a locally riskless investment over an infinitesimal interval beginning at t. The standard short-end consistency condition is

$$\boxed{x_t(0) = r_t.} \tag{1.10}$$

For a smooth discount curve, this agrees with the limit of short-horizon yields. The short rate is a boundary value of the curve, but one boundary value does not in general determine all other maturities.

1.4 Two maturity coordinates, one evolving market

This is our first important geometric distinction. The *fixed-maturity* representation uses the calendar maturity date T; the *rolling-maturity* representation uses $\tau = T - t$.

Let $f(t, T)$ denote the instantaneous forward rate at calendar time t for the fixed payment date T, and let

$$x_t(\tau) = f(t, t + \tau). \tag{1.11}$$

The two representations describe the same market information. Nevertheless, their time derivatives are different. Applying the ordinary chain rule to a sufficiently smooth deterministic function,

$$\partial_t x_t(\tau)|_{\tau \text{ fixed}} = \partial_t f(t,T)|_{T=t+\tau} + \partial_T f(t,T)|_{T=t+\tau}\,. \tag{1.12}$$

Equivalently, since $\partial_T f = \partial_\tau x$ at corresponding points,

$$\partial_t f(t,T)|_{T \text{ fixed}} = \partial_t x_t(\tau)|_{\tau \text{ fixed}} - \partial_\tau x_t(\tau). \tag{1.13}$$

For a stochastic curve, these formulas describe the coordinate transport appearing in the drift when the curve is sufficiently regular; stochastic increments require Itô calculus, introduced later.

1.4.1 A simple example: movement without repricing

Imagine a hypothetical fixed-maturity forward surface

$$f(t,T) = g(T) \tag{1.14}$$

that does not change with observation time t for any fixed future date T. Then the rolling curve is

$$x_t(\tau) = g(t+\tau). \tag{1.15}$$

Although $\partial_t f(t,T) = 0$, the rolling curve satisfies

$$\partial_t x_t(\tau) = g'(t+\tau) = \partial_\tau x_t(\tau). \tag{1.16}$$

Nothing has been repriced at a fixed maturity date. Nevertheless, the graph viewed against remaining maturity shifts as the calendar advances. This is the coordinate or *roll-down* effect. It is not a risk premium and must not be mistaken for a new economic force.

> **Common error.** An equation for $\partial_t f(t,T)$ at fixed T cannot be used unchanged as an equation for $\partial_t x_t(\tau)$ at fixed τ. The maturity-transport term must be included. Many apparent conflicts between forward-curve equations disappear once the coordinates are stated explicitly.

1.5 Realised curves, forecasts and prices are different objects

We now distinguish three concepts that ordinary language often calls a "future interest rate."

The realised forward curve today is $x_t(\tau)$. It is inferred from prices observable at date t.

The future realised forward curve is $x_{t+h}(\tau)$. It will be observable at date $t+h$, but is generally uncertain today.

The physical forecast of a future curve is the conditional expectation

$$\mathbb{E}^{\mathbb{P}}[x_{t+h}(\tau) \mid \mathcal{F}_t], \tag{1.17}$$

where $\mathcal{F}_t$ represents information available at date t and $\mathbb{P}$ denotes the physical probability law governing real-world forecasts.

There is also a fourth related quantity: the *forward price today* for lending over a future interval. It is encoded in current bond prices and need not equal the physical forecast of the corresponding future short rate. Pricing includes the effects of risk compensation and, in general, nonlinear transformations of uncertain rates.

1.5.1 Why an observed forward is not automatically a forecast

Suppose the market quotes a one-year forward rate for an interval beginning in two years. It is tempting to interpret that rate as the market's physical expectation of the short rate in two years. But a traded forward rate is derived from the prices of contingent future payments, and those prices reflect how investors value risk. Even in the absence of arbitrage, a forward rate need not coincide with a real-world expected future short rate.

This distinction foreshadows our two measures: $\mathbb{P}$ for physical forecasts and $\mathbb{Q}$ for risk-neutral valuation. We shall not use $\mathbb{Q}$ formally until the pricing chapters. For now, it is enough to recognise that *expectation* and *price* are distinct operations.

1.6 Representing an entire curve using one state

1.6.1 The one-state hypothesis

A general forward curve is a function of maturity and may require many sources of randomness to describe its evolution. We shall study the selected class in which one real-valued state r_t generates the entire curve:

$$\boxed{x_t(\tau) = x(\tau, r_t).} \tag{1.18}$$

Here $x(\tau, r)$ is a deterministic *state-to-curve map*. Given a state r, it specifies a curve across all maturities $\tau \geq 0$.

The short-end condition becomes

$$x(0, r) = r. \tag{1.19}$$

That is a boundary restriction on the map, not a statement that the rest of the curve is uniquely fixed by no-arbitrage alone. The claim that the complete curve is determined by r is the substantive *one-state modelling hypothesis*.

1.6.2 State affinity is not maturity linearity

A particularly useful class of maps is

$$\boxed{x(\tau, r) = c(\tau) + d(\tau)r.} \tag{1.20}$$

This is *affine in the state* r. It does not say that x is a straight line as a function of τ. The functions c and d may have substantial nonlinear maturity dependence.

The short-end condition (1.19) requires

$$c(0) = 0, \qquad d(0) = 1. \tag{1.21}$$

For illustration, a map such as

$$x(\tau, r) = 0.03(1 - e^{-\tau}) + e^{-\tau}r \tag{1.22}$$

is affine in r and nonlinear in τ. At the short end it equals r; at long maturities it tends to 3%. This is a *geometric example only*, not a claim that the function solves the no-arbitrage pricing equations for the diffusion to be developed later.

1.6.3 The integrated map

Combining (1.7) and (1.20), define

$$A(\tau) = \int_0^\tau c(s)\,\mathrm{d}s, \qquad B(\tau) = \int_0^\tau d(s)\,\mathrm{d}s. \tag{1.23}$$

Then the integrated forward curve and the bond-price representation become

$$X(\tau, r) := \int_0^\tau x(s, r)\,\mathrm{d}s = A(\tau) + B(\tau)r, \tag{1.24}$$

$$\mathcal{P}(\tau, r) = \exp[-A(\tau) - B(\tau)r]. \tag{1.25}$$

These are mathematical consequences of the affine curve hypothesis and the definition of instantaneous forwards. They are not yet an assertion that any arbitrary A and B yield arbitrage-free prices under a chosen short-rate diffusion. Pricing consistency will constrain them later.

Notice also that

$$A(0) = B(0) = 0, \qquad B'(0) = 1. \tag{1.26}$$

The first two identities ensure a unit payment at zero horizon costs one; the third ensures that the short end of the forward curve equals r.

1.7 Two ways of forecasting a curve

Now suppose that r_{t+h} is uncertain at date t. If the complete curve is represented by (1.18), two natural procedures suggest themselves.

Procedure A: For every possible realised future state, calculate its curve $x(\tau, r_{t+h})$ and then take the physical expectation.

$$\mathbb{E}^{\mathbb{P}}[x(\tau, r_{t+h}) \mid r_t = r]\,. \tag{1.27}$$

Procedure B: First forecast the future short rate and then evaluate the curve at the forecast state.

$$x\Big(\tau, \mathbb{E}^{\mathbb{P}}[r_{t+h} \mid r_t = r]\Big)\,. \tag{1.28}$$

In general, these operations are *not* equal. For example, if $x(\tau, r)$ contains a term proportional to r^2, then the first procedure depends on $\mathbb{E}[r_{t+h}^2]$ while the second depends only on $(\mathbb{E}[r_{t+h}])^2$. Their difference involves the forecast variance.

Our paper selects the following equality as a forecast-representation condition:

$$\boxed{\mathbb{E}_r^{\mathbb{P}}[x(\tau, r_{t+h})] = x\Big(\tau, \mathbb{E}_r^{\mathbb{P}}[r_{t+h}]\Big)\,.} \tag{1.29}$$

The notation $\mathbb{E}_r^{\mathbb{P}}$ means that the initial short-rate state is r. The maturity τ in this expression is the remaining maturity of the *future realised curve*; this is not the same as asking for the expected future value of today's bond with its fixed payment date.

1.7.1 Why affinity is sufficient

If $x(\tau, r) = c(\tau) + d(\tau)r$, and the necessary expectations exist, then

$$\mathbb{E}_r^{\mathbb{P}}[x(\tau, r_{t+h})] = \mathbb{E}_r^{\mathbb{P}}[c(\tau) + d(\tau)r_{t+h}] \tag{1.30}$$

$$= c(\tau) + d(\tau)\mathbb{E}_r^{\mathbb{P}}[r_{t+h}] \tag{1.31}$$

$$= x\Big(\tau, \mathbb{E}_r^{\mathbb{P}}[r_{t+h}]\Big)\,. \tag{1.32}$$

No choice of physical drift has been required for these equalities: linearity of expectation is enough.

1.7.2 Why the converse requires hypotheses

It would be incorrect to say that (1.29) forces affinity for every imaginable random variable or every possible short-rate process. A degenerate state with no conditional variance, for example, cannot reveal curvature in x. The rigorous selection in our research uses a sufficiently nondegenerate diffusion, smoothness and equality over initial states and short forecast horizons.

The idea can be seen using a second-order expansion. If a small forecast innovation has conditional variance of order h, then a curved function of the state introduces a correction proportional to

$$\frac{1}{2}\,\partial_{rr}x(\tau,r)\,\mathrm{Var}_r(r_{t+h}). \tag{1.33}$$

Requiring the two forecasting procedures to agree for all relevant states and sufficiently small h then removes that curvature where the diffusion variance is positive. The precise generator argument will be proved in Chapter 3.

Assumption versus deduction. We *select* the expectation-commutation principle as a representation requirement. Under stated nondegeneracy and regularity conditions, we *deduce* that the curve map is affine in the state. Affinity is not claimed to follow from the existence of bond prices or from no-arbitrage alone.

1.8 Physical equilibrium is a probability law, not a motionless curve

1.8.1 Stationarity of the state distribution

Suppose the state follows a time-homogeneous stochastic process under the physical law $\mathbb{P}$. A probability distribution π_P is *invariant* if, whenever r_0 has distribution π_P, the state r_t also has distribution π_P for every $t \geq 0$.

The state itself need not be constant. Individual sample paths may fluctuate continuously while the distribution remains unchanged. This is the meaning of *physical stationarity* used in the monograph; it is not a claim about a general equilibrium of economic agents.

If the stationary mean exists, define

$$\theta_P := \mathbb{E}_{\pi_P}[r]. \tag{1.34}$$

For an affine curve, its stationary physical mean is

$$\begin{aligned} x_e^P(\tau) &:= \mathbb{E}_{\pi_P}[x(\tau,r)] && (1.35)\\ &= c(\tau) + d(\tau)\theta_P && (1.36)\\ &= x(\tau,\theta_P). && (1.37) \end{aligned}$$

The notation x_e^P denotes the *physical equilibrium forward curve*. It is a distributional average, not a curve that every realised market path must follow.

1.8.2 A subtle but important non-implication

Stationarity means that the *unconditional* mean is constant when the process begins in its invariant law:

$$\mathbb{E}_{\pi_P}[r_t] = \theta_P \quad \text{for all } t. \tag{1.38}$$

It does *not* generally imply that a process started deterministically at $r_0 = \theta_P$ obeys

$$\mathbb{E}^{\mathbb{P}}[r_h \mid r_0 = \theta_P] = \theta_P \quad \text{for every } h. \tag{1.39}$$

That conditional claim is stronger and requires additional structure. Later we will see why it matters when deciding whether to restrict the physical drift to be affine.

1.8.3 The equilibrium curve need not be increasing

Even if π_P exists, equation (1.37) imposes no universal sign on $\partial_\tau x_e^P$. A nondecreasing stationary expected forward curve is therefore an *optional geometric selection*, not a consequence of stationarity. Our research will show how this additional selection restricts parameters once the pricing equations are imposed.

1.9 The difference between physical and pricing descriptions

We can now state why no-arbitrage pricing is a separate step. A physical model describes probabilities of future short-rate paths. Bond prices describe the market's valuation of future payments. In an arbitrage-free market, suitable pricing representations can often be expressed using a risk-neutral probability law $\mathbb{Q}$ that differs from the physical law $\mathbb{P}$.

The precise construction of $\mathbb{Q}$ and the required assumptions will come later. For now, write symbolically

$$\underbrace{b_P(r)}_{\substack{\text{physical drift}\\ \text{forecast dynamics}}} \qquad \text{and} \qquad \underbrace{b_Q(r)}_{\substack{\text{risk-neutral drift}\\ \text{pricing dynamics}}} . \tag{1.40}$$

These need not be equal. The difference is connected to compensation for bearing risk. Hence a state process can be mean-reverting under its physical law while its risk-neutral drift is not mean-reverting. There is no contradiction: they answer different questions about the same underlying uncertainty.

Later we shall obtain, with a carefully stated sign convention,

$$b_Q(r) = b_P(r) + \sigma(r)\lambda(r), \tag{1.41}$$

where λ is the market price of risk and σ the short-rate diffusion coefficient. This identity is offered here as a preview, not as a result proved in Chapter 1.

The final compatibility issue is then unavoidable: if the pricing equations constrain b_Q, what must hold between physical forecasts b_P and risk compensation λ? The monograph will answer this question explicitly, including through expected excess zero-coupon bond returns.

1.10 The hierarchy: selections, deductions and unresolved quantities

We can now lay out the intended order of reasoning. A hierarchy is not necessarily a chain of logical implications. Some of its steps are independent modelling requirements applied to the same object. Keeping this distinction visible is central to the monograph.

Question	Selection or requirement	What it can establish
Representing forecasts	Expectation commutation; regular nondegenerate state dynamics	State affinity under hypotheses
Valuing bonds	No-arbitrage under a pricing law	Restrictions on A, B, b_Q, σ
Describing long-run observations	Invariant physical law π_P	An equilibrium mean curve
Linking forecasts and returns	Consistent physical/pricing measure change	Risk compensation versus b_P
Restricting equilibrium geometry	Optional monotonicity of x_e^P	Additional parameter inequalities

A central principle of compatibility. A restrictive pricing structure may coexist with many physical stationary dynamics. The role of risk compensation is to reconcile the two descriptions. Compatibility can therefore identify a *family* of admissible models without selecting a unique physical law.

This chapter makes no claim that the observed market is exactly one-state, that real forward curves are everywhere positive, or that an interior volatility hump has been empirically established. These are separate questions. Our later classification is conditional on the precise specifications introduced in the mathematical chapters.

1.11 What Chapter 1 has established—and what remains free

We have defined zero-coupon prices, yields and instantaneous forwards; derived the yield-as-average identity; distinguished fixed-payment-date from rolling-maturity derivatives; formulated the one-state map; and explained the expectation-commutation selection and physical equilibrium interpretation.

We have established as identities:

$$\begin{aligned} y_t(\tau) &= \frac{1}{\tau}\int_0^\tau x_t(s)\,\mathrm{d}s, \\ x_t(\tau) &= y_t(\tau) + \tau\partial_\tau y_t(\tau), \\ \mathcal{P}(\tau, r) &= e^{-A(\tau)-B(\tau)r} \quad \text{if } x(\tau, r) = c(\tau) + d(\tau)r, \\ x_e^P(\tau) &= x(\tau, \theta_P) \quad \text{if affinity and finite stationary mean hold.} \end{aligned}$$

We have selected for later investigation: a one-state representation and physical forecast commutation. Neither has been derived from no-arbitrage.

We have not yet determined: the physical or pricing short-rate drift, the volatility function, the Riccati equations, the market price of risk, the existence of an invariant density, or the geometry of an equilibrium forward curve under pricing restrictions. These are deliberately reserved for subsequent chapters.

Bridge to Chapter 2

We next introduce the stochastic tools needed to study a moving state r_t: Brownian motion, diffusion equations, conditional expectations and Itô's formula. The goal is not to learn stochastic calculus for its own sake. It is to understand the operators that take us from assumptions about short-rate motion to testable equations for forecasts, bond prices and long-run probability distributions.

Chapter 2

Stochastic Dynamics, Itô Calculus and the Generator

Chapter 1 treated the interest-rate term structure as a family of bond prices and forward rates that changes with calendar time. To describe that change using a stochastic short-rate state, we need only a small set of mathematical tools, but we need to understand them precisely. This chapter develops Brownian motion, stochastic differential equations, quadratic variation, Itô's formula, conditional expectations and the infinitesimal generator. Every calculation is tied to the same short-rate state. We do not yet impose risk-neutral pricing, physical stationarity or CIR dynamics. The purpose is to establish the machinery from which those later restrictions will be derived.

2.1 From a changing rate to a random process

Chapter 1 distinguished a curve observed at date t from the random curve that may be observed at a later date. If the entire curve is represented by one state r_t, forecasting the curve requires an account of how r_t changes. A deterministic differential equation such as $\mathrm{d}r_t/\,\mathrm{d}t = b(r_t)$ would permit only one future trajectory once r_0 were specified. Actual rates are exposed to unexpected information. We therefore add a random innovation.

Throughout the mathematical foundations we write $\mathbb{P}$ for the *physical* probability law: the law under which forecasts concern what may actually happen. Under $\mathbb{P}$ a time-homogeneous one-dimensional short-rate diffusion is represented as

$$\mathrm{d}r_t = b_P(r_t)\ \mathrm{d}t + \sigma(r_t)\ \mathrm{d}W_t^P. \tag{2.1}$$

The function b_P is the physical *drift*; σ is the *diffusion coefficient*; W^P is Brownian motion. The drift and diffusion are specified functions, not yet inferred from bond prices. No claim of equilibrium is contained in (2.1).

An expression such as $\mathrm{d}W_t/\,\mathrm{d}t$ must not be treated as an ordinary random velocity. Brownian paths are almost surely nowhere differentiable. Equation (2.1) is shorthand for an *integral* equation,

$$r_t = r_0 + \int_0^t b_P(r_s)\ \mathrm{d}s + \int_0^t \sigma(r_s)\ \mathrm{d}W_s^P. \tag{2.2}$$

The first integral is an ordinary time integral; the second is an Itô stochastic integral, defined through limits of sums with the integrand evaluated at the beginning of each small interval. This convention is consequential: different conventions require different drift adjustments.

Interpretation. The drift describes the conditional tendency over a very short horizon. The diffusion describes the magnitude of unpredictable changes. Neither says by itself whether the process admits a long-run stationary distribution.

2.2 Brownian motion: three properties that matter

A standard Brownian motion W_t starts at zero and has continuous paths. Over $h > 0$ its increment $W_{t+h} - W_t$ is independent of information available up to t and is normally distributed with mean zero and variance h. For the information $\mathcal{F}_t$ available at t,

$$\mathbb{E}[W_{t+h} - W_t \mid \mathcal{F}_t] = 0, \tag{2.3}$$
$$\mathbb{E}[(W_{t+h} - W_t)^2 \mid \mathcal{F}_t] = h. \tag{2.4}$$

The standard deviation of a Brownian increment is therefore $\sqrt{h}$, not h. This simple scaling is responsible for the second-derivative term in Itô's formula.

For a small step h, the diffusion model suggests the approximation

$$r_{t+h} - r_t \approx b_P(r_t)h + \sigma(r_t)\sqrt{h}\, Z, \qquad Z \sim N(0,1), \tag{2.5}$$

conditionally on the current state. This is the Euler–Maruyama step. It is an approximation to the process, not an exact finite-horizon distribution except in special cases. Its conditional mean and variance satisfy, under regularity,

$$\mathbb{E}[r_{t+h} - r_t \mid r_t = r] = b_P(r)h + o(h), \tag{2.6}$$
$$\mathrm{Var}(r_{t+h} - r_t \mid r_t = r) = \sigma^2(r)h + o(h). \tag{2.7}$$

These equations give operational meanings to b_P and σ^2.

2.2.1 Why ordinary calculus is not enough

For a smooth deterministic path, $(\Delta r)^2$ is of smaller order than Δt whenever Δr is of order Δt. But the stochastic part of (2.5) is of order $\sqrt{\Delta t}$. Its square is therefore of order Δt:

$$\left(\sigma(r)\sqrt{\Delta t}\, Z\right)^2 = \sigma^2(r)\Delta t\, Z^2.$$

As the time partition is refined, the squares of Brownian increments accumulate rather than vanish. The resulting *quadratic variation* is

$$[W]_t = t, \qquad [r]_t = \int_0^t \sigma^2(r_s)\, \mathrm{d}s, \tag{2.8}$$

for the diffusion in (2.1). In the symbolic notation used in stochastic calculus this is often recorded as $(\mathrm{d}W_t)^2 = \mathrm{d}t$, but that notation is a compact mnemonic for quadratic variation, not a literal identity between infinitesimals.

2.3 Deriving Itô's formula from a Taylor expansion

Consider a twice differentiable function $\phi(r)$ and a small state change Δr. An ordinary Taylor expansion gives

$$\Delta\phi \approx \phi'(r)\Delta r + \frac{1}{2}\phi''(r)(\Delta r)^2. \tag{2.9}$$

Insert $\Delta r \approx b_P(r)h + \sigma(r)\sqrt{h}Z$. To first order in h, the terms have the following sizes:

$$\phi'(r)\Delta r \approx \phi'(r)b_P(r)h + \phi'(r)\sigma(r)\sqrt{h}Z,$$
$$\frac{1}{2}\phi''(r)(\Delta r)^2 \approx \frac{1}{2}\phi''(r)\sigma^2(r)hZ^2.$$

Terms of order $h^{3/2}$ and h^2 disappear in the limiting calculation. Since the squared stochastic increments accumulate to quadratic variation, the Z^2 factor contributes its limiting quadratic-variation coefficient 1. This motivates the Itô formula

$$\boxed{\mathrm{d}\phi(r_t) = \left[b_P(r_t)\phi'(r_t) + \frac{1}{2}\sigma^2(r_t)\phi''(r_t)\right]\,\mathrm{d}t + \sigma(r_t)\phi'(r_t)\,\mathrm{d}W_t^P.} \tag{2.10}$$

For a function also depending on time, $F(t, r)$, the formula adds $\partial_t F$ to the drift:

$$\mathrm{d}F(t, r_t) = \left[F_t + b_P F_r + \frac{1}{2}\sigma^2 F_{rr}\right]\,\mathrm{d}t + \sigma F_r\,\mathrm{d}W_t^P, \tag{2.11}$$

where the derivatives are evaluated at (t, r_t). Standard hypotheses include a suitable diffusion solution and $F \in C^{1,2}$; the identity may also be used in more general settings with additional care.

The crucial correction. Ordinary chain-rule reasoning would give only $b_P\phi'$. Stochastic calculus adds $\frac{1}{2}\sigma^2\phi''$. That correction is the mathematical source of both the generator's second derivative and the diffusion-variance term in the later Riccati pricing equation.

2.3.1 Worked example: the square of the short rate

Take $\phi(r) = r^2$, so $\phi'(r) = 2r$ and $\phi''(r) = 2$. Itô's formula yields

$$\mathrm{d}(r_t^2) = \left[2r_t b_P(r_t) + \sigma^2(r_t)\right]\,\mathrm{d}t + 2r_t\sigma(r_t)\,\mathrm{d}W_t^P. \tag{2.12}$$

The extra σ^2 term means that the evolution of the second moment cannot generally be deduced by squaring the expected evolution of the first moment. Taking expectations (when integrability permits) gives

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[r_t^2] = \mathbb{E}\left[2r_t b_P(r_t) + \sigma^2(r_t)\right]. \tag{2.13}$$

The first-moment equation is simpler:

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[r_t] = \mathbb{E}[b_P(r_t)]. \tag{2.14}$$

In general $\mathbb{E}[b_P(r_t)] \neq b_P(\mathbb{E}[r_t])$. Equality requires an affine drift or some other special property. This distinction will matter when we discuss forecast consistency and mean dynamics.

2.3.2 Worked example: an exponential function

Let $\phi(r) = e^{-Br}$ for a fixed constant B. Then $\phi' = -B\phi$ and $\phi'' = B^2\phi$. Hence

$$\mathrm{d}e^{-Br_t} = e^{-Br_t}\left[-Bb_P(r_t) + \frac{1}{2}B^2\sigma^2(r_t)\right]\,\mathrm{d}t - Be^{-Br_t}\sigma(r_t)\,\mathrm{d}W_t^P. \tag{2.15}$$

Notice that the diffusion contributes a *positive* $B^2\sigma^2/2$ to the drift of the exponential, even though its stochastic increment has mean zero. Later, substituting an exponential-affine *bond price* into the risk-neutral pricing equation will produce this same term, with b_Q in place of b_P.

2.4 The infinitesimal generator

The notation for the recurring drift term is important enough to deserve its own definition.

$$\boxed{\mathcal{L}_P\phi(r) := b_P(r)\phi'(r) + \frac{1}{2}\sigma^2(r)\phi''(r).} \tag{2.16}$$

The operator $\mathcal{L}_P$ is the *physical infinitesimal generator* of the diffusion. It is linear as an operator acting on test functions ϕ, even when b_P and σ are nonlinear functions of r.

A useful way to think about the generator is as an infinitesimal forecasting machine. For a suitable smooth function ϕ,

$$\mathcal{L}_P\phi(r) = \lim_{h\downarrow 0} \frac{\mathbb{E}_r^P[\phi(r_h)] - \phi(r)}{h}, \tag{2.17}$$

where $\mathbb{E}_r^P$ means conditional expectation given $r_0 = r$. For $\phi(r) = r$ and $\phi(r) = r^2$ we obtain

$$\mathcal{L}_P r = b_P(r), \tag{2.18}$$

$$\mathcal{L}_P r^2 = 2rb_P(r) + \sigma^2(r). \tag{2.19}$$

The generator acts on an *observable function* of the state. It is not itself a probability density, and it is not the same thing as the Fokker–Planck operator that acts on densities.

2.4.1 Why conditional expectations produce the generator

From Itô's formula over $[0, h]$,

$$\phi(r_h) - \phi(r) = \int_0^h \mathcal{L}_P\phi(r_s)\ \mathrm{d}s + \int_0^h \sigma(r_s)\phi'(r_s)\ \mathrm{d}W_s^P.$$

Under the usual integrability conditions, the Itô integral has expectation zero. Therefore,

$$\mathbb{E}_r^P[\phi(r_h)] - \phi(r) = \mathbb{E}_r^P\left[\int_0^h \mathcal{L}_P\phi(r_s)\ \mathrm{d}s\right]. \tag{2.20}$$

Divide by h and let h tend to zero. If continuity and domination permit the limit to pass through expectation, we obtain (2.17). Equation (2.20) is a form of *Dynkin's formula*. It gives a precise meaning to the frequently used phrase "differentiate the expected forecast at the initial horizon."

2.5 From a single state to expected future curves

Chapter 1 introduced a state-to-forward-curve mapping $x(\tau, r)$. At date t the curve is $x(\tau, r_t)$. Under physical forecasting, its expected value after horizon h is

$$m_P(h, \tau; r) := \mathbb{E}_r^P[x(\tau, r_h)]. \tag{2.21}$$

The maturity coordinate τ in this expression is a chosen *remaining maturity at the future observation date*; it is not the fixed original bond maturity T. This preserves the distinction made in Chapter 1.

Treat τ as fixed and apply the generator to the state variable r. Provided the relevant derivatives and expectations exist,

$$\boxed{\partial_h m_P(h, \tau; r)|_{h=0} = \mathcal{L}_P x(\tau, r) = b_P(r)\partial_r x(\tau, r) + \frac{1}{2}\sigma^2(r)\partial_{rr} x(\tau, r).} \tag{2.22}$$

This identity is the first bridge from the stochastic short-rate model to forecasts of the complete curve.

2.5.1 The special simplification under state affinity

Suppose, as an explicitly selected model class rather than a conclusion of this chapter,

$$x(\tau, r) = A_x(\tau) + C_x(\tau)r. \tag{2.23}$$

The second state derivative vanishes: $\partial_{rr}x = 0$. The generator identity simplifies to

$$\mathcal{L}_P x(\tau, r) = C_x(\tau) b_P(r). \tag{2.24}$$

More strongly, at any horizon with a finite first moment,

$$\mathbb{E}_r^P[x(\tau, r_h)] = A_x(\tau) + C_x(\tau)\mathbb{E}_r^P[r_h] = x\Big(\tau, \mathbb{E}_r^P[r_h]\Big). \tag{2.25}$$

Thus affine state dependence is sufficient for the expectation-commutation relationship discussed in Chapter 1. This calculation does not assert that commutation uniquely implies affinity without further assumptions on the family of short-rate transition laws. That more delicate converse belongs in Chapter 3.

A revealing distinction. For an affine curve, taking the expectation of the curve is straightforward even when $b_P(r)$ is nonlinear. But forecasting the *mean short rate* may remain difficult, because $\partial_h \mathbb{E}[r_h] = \mathbb{E}[b_P(r_h)]$ generally does not close as an ordinary differential equation for $\mathbb{E}[r_h]$ alone.

2.6 An illustrative restoring diffusion: no CIR assumption required

For orientation, consider the linear restoring physical drift

$$\mathrm{d}r_t = \kappa(\theta - r_t)\,\mathrm{d}t + \eta\,\mathrm{d}W_t^P, \qquad \kappa > 0,\ \eta > 0. \tag{2.26}$$

This is the Ornstein–Uhlenbeck model. It is mathematically transparent but, because its state is Gaussian, it does *not* enforce positivity of interest rates. It is an example for learning, not our preferred positive short-rate specification.

Take expectations conditional on $r_0 = r$. Since the drift is affine,

$$\frac{\mathrm{d}}{\mathrm{d}h}\mathbb{E}_r^P[r_h] = \kappa(\theta - \mathbb{E}_r^P[r_h]). \tag{2.27}$$

Solving this ordinary differential equation gives

$$\mathbb{E}_r^P[r_h] = \theta + (r - \theta)e^{-\kappa h}. \tag{2.28}$$

Using the explicit stochastic solution or applying Itô calculus, the conditional variance is

$$\mathrm{Var}_r^P(r_h) = \frac{\eta^2}{2\kappa}(1 - e^{-2\kappa h}). \tag{2.29}$$

Both equations illustrate the generator's task: the drift determines the first conditional moment, while the diffusion contributes directly to the second. In this particular model, the moment equations close because the drift is affine and the variance coefficient is constant.

Now contrast the positive square-root volatility relevant to our research:

$$\sigma(r) = \sqrt{\beta r}, \qquad r \geq 0, \quad \beta > 0. \tag{2.30}$$

The conditional instantaneous variance rate is βr. This is a statement about local randomness; it does not, on its own, determine the physical drift. We may combine (2.30) with an admissible nonlinear restoring drift b_P and obtain quite different physical stationary laws. The requirements at the boundary $r = 0$ must be checked separately. We will do so in the equilibrium chapters.

2.7 The generator under a change of probability measure

So far everything has been under $\mathbb{P}$. Later, no-arbitrage bond pricing will introduce a second probability measure, $\mathbb{Q}$. In the compatibility paper our convention is

$$\mathrm{d}W_t^Q = \mathrm{d}W_t^P - \lambda(r_t)\,\mathrm{d}t, \qquad b_Q(r) = b_P(r) + \sigma(r)\lambda(r). \tag{2.31}$$

This chapter does *not* prove that such a measure change exists: an equivalent change requires appropriate integrability conditions on the market-price-of-risk process $\lambda(r_t)$, as will be discussed later. For now, the equation is a preview of the bookkeeping.

When it is valid, the new generator is

$$\mathcal{L}_Q\phi = b_Q\phi' + \frac{1}{2}\sigma^2\phi'', \tag{2.32}$$

so that

$$\boxed{(\mathcal{L}_Q - \mathcal{L}_P)\phi(r) = \sigma(r)\lambda(r)\phi'(r).} \tag{2.33}$$

The second-derivative coefficients agree; only the drift changes. The difference of generators is therefore a *first-order* operator. Applied to $\phi(r) = r$, it gives $b_Q - b_P = \sigma\lambda$. Applied to the affine forward map $x(\tau, r) = A_x + C_x r$, it gives

$$(\mathcal{L}_Q - \mathcal{L}_P)x(\tau, r) = C_x(\tau)\sigma(r)\lambda(r). \tag{2.34}$$

These are the foundations of the later connection between physical forecasts, risk-neutral drift and bond risk compensation. We are not yet equating a forward-rate generator difference with a bond excess-return rate; that further connection requires the bond-price dynamics and will be derived explicitly in the pricing chapters.

2.8 What Chapter 2 has established—and what remains free

We now have a precise description of what it means for a short-rate state to evolve stochastically. Brownian increments scale as the square root of elapsed time; their quadratic variation produces the extra second derivative in Itô's formula. The generator packages the local drift of any sufficiently regular function of the state. It also determines the initial derivative of a physical conditional forecast.

For a state-affine forward curve, expected future curves commute with taking the expected state, but this does not by itself determine the physical mean dynamics. A fixed choice of diffusion coefficient likewise does not identify a physical drift or a stationary law. Under an admissible change of measure, the physical and risk-neutral generators differ in their first-derivative coefficients while sharing the same local variance coefficient.

Established here: the stochastic-calculus identities and their conditional-forecast interpretation. **Not established here:** that the short rate is positive; that a physical invariant density exists; that x must be affine; that $\mathbb{Q}$ exists or prices bonds; or that the physical drift is CIR. Each requires separate assumptions or arguments.

Bridge to Chapter 3

We have constructed the generator as the mathematical connection between stochastic state dynamics and short-horizon expected changes. The next chapter will study *forecast consistency*

more deeply: what it means to represent an entire future curve by a single state, why affinity is sufficient for expectation commutation, which extra hypotheses are required for a converse, and how one avoids confusing a forecast under $\mathbb{P}$ with an arbitrage-free price under $\mathbb{Q}$.

Chapter 3

Forecast Consistency and the Affine State-to-Curve Map

Chapters 1 and 2 introduced the evolving term structure and the infinitesimal generator of a stochastic short rate. We now ask whether forecasting the entire future curve by averaging possible future curves should agree with evaluating the curve at the expected future short rate. These operations do not ordinarily commute. This chapter derives the exact expectation-commutation condition, explains why an affine state-to-curve map satisfies it, and proves a converse under explicit differentiability and nondegeneracy assumptions. We then separate affinity of the curve from affinity of the physical drift, derive the implications for stationary expected curves, and explain what this selection does not establish about pricing or the shape of the maturity curve.

3.1 The problem: two ways to forecast one curve

At date t, let $r_t = r$ be the observed short rate. Under the physical law $\mathbb{P}$, suppose that its future evolution over a horizon $h \geq 0$ is described by a time-homogeneous Markov diffusion

$$\mathrm{d}r_s = b_P(r_s)\ \mathrm{d}s + \sigma(r_s)\ \mathrm{d}W_s^P. \tag{3.1}$$

Write $x(\tau, r)$ for the instantaneous-forward curve generated by state r at remaining maturity $\tau \geq 0$. The short-end boundary is

$$x(0, r) = r. \tag{3.2}$$

Here h and τ perform different jobs. The horizon h is how far we look into the future from today; τ is the time remaining until the future payment, measured *at that future observation date*. We hold τ fixed when comparing future curves. This is the rolling-maturity viewpoint established in Chapter 1.

There are two natural forecasting procedures. The first averages the full curve across possible future states:

$$M(h, \tau; r) := \mathbb{E}_r^P\left[x(\tau, r_h)\right]. \tag{3.3}$$

The second predicts the future short rate first and then inserts that expected state into today's state-to-curve map:

$$\widehat{M}(h, \tau; r) := x(\tau, m(h, r)), \qquad m(h, r) := \mathbb{E}_r^P[r_h]. \tag{3.4}$$

These procedures may look interchangeable. In general they are not: the expectation of a nonlinear function differs from that function evaluated at the expectation of its argument.

The forecasting selection. The exact closure condition studied in this monograph is

$$\boxed{\mathbb{E}_r^P[x(\tau, r_h)] = x(\tau, \mathbb{E}_r^P[r_h]).} \tag{3.5}$$

We adopt this equality as an explicit modelling selection when it is required for every state, maturity and sufficiently short forecast horizon. It is *not* an identity enjoyed by arbitrary forward curves.

3.2 Why a nonlinear curve generally fails the test

Before discussing interest rates, consider a random variable R with finite variance. If $g(R) = R^2$, then

$$\mathbb{E}[g(R)] = \mathbb{E}[R^2], \tag{3.6}$$
$$g(\mathbb{E}[R]) = (\mathbb{E}[R])^2. \tag{3.7}$$

Their difference is exactly

$$\mathbb{E}[R^2] - (\mathbb{E}[R])^2 = \mathrm{Var}(R). \tag{3.8}$$

Unless R is deterministic, the two procedures disagree. This familiar identity is the simplest illustration of what nonlinearity does to averages.

For a concrete forward-curve example, consider the *illustrative*, not necessarily arbitrage-free, map

$$x(\tau, r) = r + \eta\tau r^2, \qquad \eta > 0. \tag{3.9}$$

It respects the short boundary $x(0, r) = r$. At forecast horizon h,

$$\mathbb{E}_r^P[x(\tau, r_h)] = m(h, r) + \eta\tau\mathbb{E}_r^P[r_h^2], \tag{3.10}$$
$$x(\tau, m(h, r)) = m(h, r) + \eta\tau m(h, r)^2. \tag{3.11}$$

Subtracting gives

$$M(h, \tau; r) - \widehat{M}(h, \tau; r) = \eta\tau\,\mathrm{Var}_r^P(r_h). \tag{3.12}$$

For positive τ and nonzero conditional variance, closure fails. No pricing argument was needed to reach this conclusion.

The sign of the discrepancy reflects curvature. A convex function usually has $\mathbb{E}[g(R)] \geq g(\mathbb{E}[R])$ (Jensen's inequality), while a concave function reverses the inequality. But our selection is stronger than an inequality: it requires *equality* across a whole family of conditional laws. That is what makes it restrictive.

3.3 Why an affine map works for every physical forecast law

Consider a state-affine forward curve

$$x(\tau, r) = c(\tau) + d(\tau)r. \tag{3.13}$$

Affine means a constant plus a coefficient multiplied by the state. The functions c and d may be highly nonlinear functions of *maturity*. In particular, affinity does not mean that the forward curve is a straight line as maturity increases.

Take the physical conditional expectation:

$$\mathbb{E}_r^P[x(\tau, r_h)] = \mathbb{E}_r^P[c(\tau) + d(\tau)r_h] \tag{3.14}$$
$$= c(\tau) + d(\tau)\mathbb{E}_r^P[r_h] \tag{3.15}$$
$$= x(\tau, \mathbb{E}_r^P[r_h]). \tag{3.16}$$

This is ordinary linearity of expectation. It holds regardless of whether b_P is linear, nonlinear, mean-reverting or non-restoring, provided the conditional first moment exists.

Proposition 3.1 (Affinity is sufficient for forecast closure)**.** *If $x(\tau, r) = c(\tau) + d(\tau)r$ and $\mathbb{E}_r^P|r_h| < \infty$, then the expectation-commutation identity* (3.5) *holds for every state r, maturity τ and horizon h for which the state process is defined.*

Proof. The three displayed equalities above prove the result directly. They require no SDE calculation and no stationarity assumption. □

The short boundary (3.2) implies two conditions:

$$c(0) = 0, \qquad d(0) = 1. \tag{3.17}$$

These conditions do not specify the maturity dependence of c or d. That dependence will be restricted later by arbitrage-free pricing.

3.4 The converse: when forecast closure forces affinity

The more interesting question runs in the other direction. Suppose we did *not* start by assuming that $x(\tau, r)$ is affine. Would exact closure force it to be affine? Under the appropriate hypotheses, yes.

We use the physical generator developed in Chapter 2:

$$\mathcal{L}_P\phi(r) = b_P(r)\phi'(r) + \frac{1}{2}\sigma^2(r)\phi''(r). \tag{3.18}$$

Fix a maturity τ and an interior starting state r. Consider a short forecasting horizon h. Under sufficient regularity to apply the generator expansion,

$$\mathbb{E}_r^P[x(\tau, r_h)] = x(\tau, r) + h\mathcal{L}_P x(\tau, r) + o(h). \tag{3.19}$$

Inserting (3.18) yields

$$\mathbb{E}_r^P[x(\tau, r_h)] = x(\tau, r) + h\left[b_P(r)\partial_r x(\tau, r) + \frac{1}{2}\sigma^2(r)\partial_{rr}x(\tau, r)\right] + o(h). \tag{3.20}$$

The forecast of the short rate itself satisfies

$$m(h, r) = \mathbb{E}_r^P[r_h] = r + b_P(r)h + o(h), \tag{3.21}$$

because $\mathcal{L}_P r = b_P(r)$. We now take an ordinary first-order Taylor expansion of the function $x(\tau, \cdot)$ around r:

$$x(\tau, m(h, r)) = x(\tau, r) + \partial_r x(\tau, r)[m(h, r) - r] + o(h) \tag{3.22}$$
$$= x(\tau, r) + hb_P(r)\partial_r x(\tau, r) + o(h). \tag{3.23}$$

Finally subtract (3.23) from (3.20). The physical drift term cancels, leaving

$$\boxed{\mathbb{E}_r^P[x(\tau, r_h)] - x(\tau, \mathbb{E}_r^P[r_h]) = \frac{1}{2}\sigma^2(r)\partial_{rr}x(\tau, r)h + o(h).} \tag{3.24}$$

This calculation is worth dwelling on. The first-order effect of the expected change in the short rate appears in both forecasts and cancels. What remains is the effect of *uncertainty interacting with curvature*. It is the same mathematical mechanism that produced the Itô correction in Chapter 2.

Proposition 3.2 (Local nondegenerate closure selects state affinity)**.** *Let the short-rate state range over a connected open interval I. For every fixed maturity τ, suppose $x(\tau, \cdot)$ is twice continuously differentiable on I, the physical diffusion variance satisfies $\sigma^2(r) > 0$ for every $r \in I$, and the short-horizon generator and mean expansions used in* (3.24) *are valid. If* (3.5) *holds for every starting state $r \in I$, every maturity τ, and all sufficiently small positive horizons h, then*

$$x(\tau, r) = c(\tau) + d(\tau)r \quad (r \in I). \tag{3.25}$$

Proof. Under closure, the left-hand side of (3.24) is zero. Divide (3.24) by h and let $h \downarrow 0$ to obtain

$$0 = \frac{1}{2}\sigma^2(r)\partial_{rr}x(\tau, r).$$

Since $\sigma^2(r) > 0$ on I, we have $\partial_{rr}x(\tau, r) = 0$ there. Integrating twice with respect to r on the connected interval gives (3.25). This proves affinity in the state, for each fixed maturity. □

Scope of the proposition. The converse is not a statement that every one-state term-structure model is affine. It depends on an *exact forecasting selection*, its validity across states and small horizons, a sufficiently smooth curve map and nonzero interior variance. If these conditions are relaxed, the conclusion need not follow. In particular, equality observed at one state or one horizon is not enough.

3.5 Examples showing why the hypotheses matter

3.5.1 A deterministic process

If $\sigma(r) \equiv 0$, then conditional on $r_0 = r$ the future state is deterministic (under ordinary uniqueness conditions). Therefore

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

for *any* curve map, including nonlinear ones. With no random dispersion, expectation-commutation gives no information about curvature. This explains why nondegenerate diffusion is essential to Proposition 3.2.

3.5.2 Equality at only one maturity

Every valid curve map satisfies $x(0, r) = r$, so closure at $\tau = 0$ is automatic. It says nothing about the shape of $x(\tau, r)$ when $\tau > 0$. The assumption must therefore apply across maturities, not just at the short boundary.

3.5.3 Approximate rather than exact closure

If a nonlinear curve has small curvature or the short rate has little conditional variance, the difference in (3.24) can be small. Empirical approximate closure might therefore be a useful approximation even when exact affinity is false. The theorem concerns exact mathematical closure, not a threshold for statistical significance.

3.6 Affinity of the curve does not impose affine physical drift

This is a crucial separation in our framework. Suppose the curve is affine, $x(\tau, r) = c(\tau) + d(\tau)r$, and the physical short rate follows the generic diffusion (3.1). Itô's formula gives, at *fixed rolling maturity* τ,

$$\mathrm{d}x(\tau, r_t) = d(\tau)b_P(r_t)\ \mathrm{d}t + d(\tau)\sigma(r_t)\ \mathrm{d}W_t^P. \tag{3.26}$$

There is no second-derivative Itô correction here because $\partial_{rr}x = 0$. But $d(\tau)b_P(r)$ can still be a nonlinear function of r. Nothing forces b_P itself to be affine.

For example, consider a possible nonlinear restoring physical drift on the positive state domain,

$$b_P(r) = m - kr - \delta r^2, \qquad k, \delta > 0, \tag{3.27}$$

with a square-root diffusion coefficient $\sigma(r) = \sqrt{\beta r}$ and appropriate positive-boundary assumptions. The forward map $c(\tau) + d(\tau)r$ satisfies expectation-commutation, even though the conditional mean obeys

$$\frac{\mathrm{d}}{\mathrm{d}h}\mathbb{E}_r^P[r_h] = \mathbb{E}_r^P[b_P(r_h)] \tag{3.28}$$

$$= m - k\mathbb{E}_r^P[r_h] - \delta\mathbb{E}_r^P[r_h^2]. \tag{3.29}$$

The equation for the mean now depends on the second moment. It is not a closed linear ordinary differential equation for the mean alone. Nevertheless, the curve-averaging identity remains exact.

The essential distinction. Affinity of $x(\tau, r)$ is a property of the *mapping from a state to a curve*. Affinity of $b_P(r)$ is a property of the *physical evolution of that state*. The first does not imply the second.

3.7 A stronger forecast condition would select mean-reverting affine drift

It is instructive to examine a condition that we deliberately do *not* impose. Suppose that for constants θ and $\kappa > 0$, the conditional short-rate mean satisfies

$$\mathbb{E}_r^P[r_h] = \theta + (r - \theta)e^{-\kappa h} \quad \text{for every state } r \text{ and horizon } h \geq 0. \tag{3.30}$$

Differentiating at $h = 0$ gives

$$b_P(r) = \frac{\partial}{\partial h}\mathbb{E}_r^P[r_h]\bigg|_{h=0} = -\kappa(r - \theta) = \kappa(\theta - r). \tag{3.31}$$

Thus exact exponential mean relaxation across all initial states selects an affine physical drift. This is stronger than the expectation-commutation condition on the forward-curve map.

Conversely, for a nonexplosive diffusion with $b_P(r) = \kappa(\theta - r)$ and the required integrability, taking expectations gives

$$\frac{\mathrm{d}m(h,r)}{\mathrm{d}h} = \kappa(\theta - m(h,r)), \qquad m(0,r) = r,$$

whose solution is (3.30). This is why an affine mean-reverting drift gives simple conditional forecasts. But forecasting *closure of the curve* never required these simple forecasts.

The terms *stationary mean* and *conditional mean* must not be conflated. Even if a diffusion has stationary mean θ_P, starting it from $r_0 = \theta_P$ does not generally imply $\mathbb{E}^P_{\theta_P}[r_h] = \theta_P$ for all h. That stronger property holds in special drift classes, not as a consequence of stationarity alone.

3.8 Physical equilibrium as a separate averaging operation

Suppose that the physical diffusion admits an invariant probability law π_P with finite mean

$$\theta_P = \int_I r\, \pi_P(\,\mathrm{d}r). \tag{3.32}$$

If r_0 is drawn from this law, then r_h has the same law for every h. The *stationary expected forward curve* is defined by

$$x_e^P(\tau) = \int_I x(\tau, r)\, \pi_P(\,\mathrm{d}r). \tag{3.33}$$

For an affine curve, this reduces immediately to

$$\boxed{x_e^P(\tau) = c(\tau) + d(\tau)\theta_P = x(\tau, \theta_P).} \tag{3.34}$$

This equality describes an *unconditional stationary average*. It does not mean that an observed curve at state $r = \theta_P$ will remain unchanged as time evolves. Individual paths fluctuate, and even their conditional mean need not remain at θ_P when started there.

Subject to integrability and the zero-net-flux conditions appropriate to the stationary diffusion, invariance also gives

$$\int_I b_P(r)\, \pi_P(\,\mathrm{d}r) = 0. \tag{3.35}$$

This is a condition on the *average* physical drift. It does not require $b_P(\theta_P) = 0$. For the nonlinear example (3.27), stationarity would imply

$$0 = m - k\theta_P - \delta\mathbb{E}_{\pi_P}[r^2], \tag{3.36}$$

which depends on more than the stationary mean.

3.9 Forward curves, yields and bond prices under state affinity

The affine forward map has a useful geometric consequence. Define the integrated forward rate by

$$X(\tau, r) = \int_0^\tau x(u, r)\ \mathrm{d}u. \tag{3.37}$$

Integrating $x(\tau, r) = c(\tau) + d(\tau)r$ gives

$$X(\tau, r) = A(\tau) + B(\tau)r, \quad A(\tau) = \int_0^\tau c(u)\ \mathrm{d}u, \quad B(\tau) = \int_0^\tau d(u)\ \mathrm{d}u. \tag{3.38}$$

The zero-coupon bond price consistent with the *definition* of the forward rate is therefore

$$\boxed{P(\tau, r) = \exp[-A(\tau) - B(\tau)r].} \tag{3.39}$$

Differentiating in maturity recovers the forward curve:

$$-\partial_\tau \log P(\tau, r) = A'(\tau) + B'(\tau)r = c(\tau) + d(\tau)r. \tag{3.40}$$

The short-end conditions give

$$A(0) = B(0) = 0, \quad A'(0) = 0, \quad B'(0) = 1. \tag{3.41}$$

Thus the map is *exponential-affine in the state* when expressed as a bond price. This is a representational consequence of integrating an affine forward curve, not yet a pricing theorem.

An essential separation. Equation (3.39) describes how prices and forward rates are related at a given date. It does not yet assert that an arbitrary pair A, B solves the risk-neutral pricing equation. No-arbitrage will place additional differential restrictions on A and B, derived in subsequent chapters.

The zero-coupon yield $y(\tau, r) = -\tau^{-1} \log P(\tau, r)$, for $\tau > 0$, is also affine in the state:

$$y(\tau, r) = \frac{A(\tau)}{\tau} + \frac{B(\tau)}{\tau} r. \tag{3.42}$$

In the first chapter we noted that yields average forward rates over maturity. Equation (3.42) shows that their state sensitivity is correspondingly an average of the forward-rate loading $d(u)$.

3.10 Physical forecasts are not risk-neutral prices

Up to this point all forward-looking expectations in the chapter have been taken under $\mathbb{P}$. They describe physical forecasts of future states and future *curves*. The no-arbitrage valuation of a bond instead involves a pricing measure $\mathbb{Q}$, where such a measure exists. In the standard short-rate setting the pricing expression is

$$P(\tau, r) = \mathbb{E}_r^Q \left[\exp\left(- \int_0^\tau r_s \, \mathrm{d}s \right) \right]. \tag{3.43}$$

This is a different expectation from $\mathbb{E}_r^P[x(\tau, r_h)]$. In particular, $x(\tau, r)$ should not generally be interpreted as $\mathbb{E}_r^P[r_\tau]$, nor as $\mathbb{E}_r^Q[r_\tau]$. The exponential discount factor, the measure and the running integral all matter.

The next chapters will connect (3.39) and (3.43) by means of the pricing generator and will show how the maturity loading $B(\tau)$ becomes constrained by a Riccati equation. The physical drift will remain a separate object unless further information, such as expected excess returns, is supplied.

3.11 What Chapter 3 has established—and what remains free

We introduced an explicit forecasting selection: the expected future forward curve may be computed either by averaging future curves or by evaluating the state-to-curve map at the expected future short rate. An affine state-to-curve map satisfies this condition for *any* physical state process with a finite conditional mean. Conversely, with a smooth map, nonzero interior diffusion variance, and exact closure across states and short horizons, the condition forces affine dependence on the state.

The result is about the *state-to-curve map*, not about the physical drift. Even when the curve is affine, the short-rate dynamics may be nonlinear, and physical equilibrium may involve a non-Gamma distribution. Exact exponential relaxation of the conditional short-rate mean is a different, stronger condition and is not assumed here.

Finally, integrating the affine forward curve yields an exponential-affine bond-price representation. That representation has not yet been shown to be arbitrage-free: pricing restrictions and the Riccati equation enter only when we impose the appropriate risk-neutral bond-pricing equation.

Where the hierarchy now stands. Forecast closure, under stated nondegeneracy and regularity conditions, selects *state affinity*. It does not select a physical drift, stationary density, market price of risk, or risk-neutral pricing dynamics. Each additional object requires its own compatibility argument.

Bridge to Chapter 4

We are ready to move from *forecasting a curve* to *pricing a bond*. The next chapter introduces the risk-neutral bond-pricing generator and derives the restrictions it imposes on $A(\tau)$, $B(\tau)$, the short-rate diffusion variance and risk-neutral drift. In that development, we will carefully separate the definitions of forward rates from no-arbitrage conditions, and the rolling-maturity coordinates from fixed-maturity bond holding returns.

Chapter 4

Arbitrage-Free Bond Pricing and the Riccati Equations

Chapter 3 selected a state-affine forward-curve representation through an explicit physical forecasting requirement. Integrating that curve produced exponential-affine bond prices, but did not establish whether those prices could coexist with arbitrage-free trading. In this chapter we introduce the risk-neutral pricing measure, derive the zero-coupon bond-pricing equation from the martingale principle and Itô's formula, and substitute the exponential-affine representation. We show exactly why two distinct nonzero bond loadings constrain the risk-neutral drift and short-rate variance to be affine, and why a separate zero-variance boundary selects square-root diffusion. We then derive the Riccati equations, initial-maturity identities and the relation between fixed-date bond returns and rolling-maturity curves. The physical drift and physical stationary law remain deliberately unidentified.

4.1 What changes when we move from forecasting to pricing?

Chapter 3 considered the physical expectation $\mathbb{E}_r^P[x(\tau, r_h)]$. This answers a forecasting question: given today's state r, what future curve should we expect under the statistical law governing realised interest rates? A bond price answers a different question: what must be paid today for a specified cash flow at a future *calendar date*? In a frictionless arbitrage-free short-rate model, with a valid equivalent risk-neutral measure $\mathbb{Q}$ and the money-market account as numeraire, a zero-coupon bond paying one unit at date T has the valuation

$$P(t,T) = \mathbb{E}_t^Q\left[\exp\left(-\int_t^T r_s \, \mathrm{d}s\right)\right]. \tag{4.1}$$

The exponential contains the *integral of the short rate along its future path.* It is not the exponential of an instantaneous forward rate. Nor, in general, may the expectation be moved through the exponential:

$$\mathbb{E}^Q[e^{-Z}] \neq e^{-\mathbb{E}^Q[Z]}.$$

The difference is a basic convexity effect; it will be visible directly in the differential calculation below. Equation (4.1) requires the usual modelling conditions ensuring that the discounted bond price is a true martingale and the expectation is finite. The existence of such a pricing measure is an assumption of this chapter, not a consequence of the forecast consistency condition.

For a time-homogeneous one-state Markov model, write the bond value as $P(\tau, r)$, where $\tau = T - t$ is remaining maturity and the current short rate is r. Its deterministic representation need not mean that the market itself is stationary: it means the model coefficients have no explicit

calendar-time dependence. We always require

$$P(0, r) = 1, \tag{4.2}$$

because a bond with zero time remaining pays one immediately.

Two uses of expectations. The physical measure $\mathbb{P}$ describes forecasts of future states. The pricing measure $\mathbb{Q}$ prices discounted cash flows. The two measures may give different conditional means for the same short-rate process. They are not interchangeable, even when the state-to-curve map is affine.

4.2 The risk-neutral short-rate generator

Assume under $\mathbb{Q}$ that the short rate follows a one-dimensional diffusion

$$\mathrm{d}r_t = b_Q(r_t)\ \mathrm{d}t + \sigma(r_t)\ \mathrm{d}W_t^Q, \qquad v(r) := \sigma(r)^2. \tag{4.3}$$

For a smooth state function ϕ, Chapter 2 gave the risk-neutral infinitesimal generator

$$\mathcal{L}_Q\phi(r) = b_Q(r)\phi'(r) + \frac{1}{2}v(r)\phi''(r). \tag{4.4}$$

To understand bond pricing, first hold the maturity date T fixed. The bond-price process is $V_t = P(T - t, r_t)$. Itô's formula, with $\tau = T - t$ and $\mathrm{d}\tau = -\,\mathrm{d}t$, gives

$$\begin{aligned}\mathrm{d}V_t = &\left[-\partial_\tau P(\tau, r_t) + b_Q(r_t)\partial_r P(\tau, r_t) + \frac{1}{2}v(r_t)\partial_{rr}P(\tau, r_t)\right]\ \mathrm{d}t \\ &+ \sigma(r_t)\partial_r P(\tau, r_t)\ \mathrm{d}W_t^Q.\end{aligned} \tag{4.5}$$

The short-rate money-market account is $M_t = \exp(\int_0^t r_s\ \mathrm{d}s)$, satisfying $\mathrm{d}M_t = r_t M_t\ \mathrm{d}t$. A traded bond's discounted value V_t/M_t must have zero local drift under $\mathbb{Q}$. Applying the product rule to $M_t^{-1}V_t$ gives a discounted drift proportional to

$$-\partial_\tau P + \mathcal{L}_Q P - rP.$$

Setting this equal to zero yields the *risk-neutral bond-pricing equation*

$$\boxed{\partial_\tau P(\tau, r) = \mathcal{L}_Q P(\tau, r) - rP(\tau, r), \qquad P(0, r) = 1.} \tag{4.6}$$

This is the backward pricing equation expressed in remaining maturity. With sufficient regularity, growth and boundary conditions it is also the Feynman–Kac equation associated with (4.1). We have derived its signs rather than relying on memory: the $-rP$ term comes from discounting, while the positive $\partial_\tau P$ appears because elapsed time reduces remaining maturity.

4.3 The exponential-affine representation: what is already known?

Forecast closure in Chapter 3 selected the state-affine forward map

$$x(\tau, r) = c(\tau) + d(\tau)r, \qquad x(0, r) = r. \tag{4.7}$$

At each date, the instantaneous forward rate is the negative maturity derivative of the logarithm of the bond price:

$$x(\tau, r) = -\partial_\tau \log P(\tau, r). \tag{4.8}$$

Integrating from maturity zero and using $P(0, r) = 1$ gives

$$-\log P(\tau, r) = \int_0^\tau x(u, r)\, \mathrm{d}u \tag{4.9}$$

$$= \underbrace{\int_0^\tau c(u)\, \mathrm{d}u}_{A(\tau)} + \underbrace{\int_0^\tau d(u)\, \mathrm{d}u}_{B(\tau)} r. \tag{4.10}$$

Therefore

$$\boxed{P(\tau, r) = \exp[-A(\tau) - B(\tau)r].} \tag{4.11}$$

The functions A and B satisfy $A(0) = B(0) = 0$. Because $c(0) = 0$ and $d(0) = 1$ follow from $x(0, r) = r$ for all states,

$$A'(0) = 0, \qquad B'(0) = 1. \tag{4.12}$$

At this stage A and B are *not* arbitrary once we impose pricing. They must satisfy (4.6); obtaining that condition is the decisive step.

4.4 Substitution into the pricing equation, without skipped steps

For $P = e^{-A-Br}$ we calculate the maturity derivative while holding the current state fixed:

$$\partial_\tau P = -(A' + B'r)P. \tag{4.13}$$

For the state derivatives, while holding maturity fixed,

$$\partial_r P = -BP, \qquad \partial_{rr} P = B^2 P. \tag{4.14}$$

The second derivative is positive because differentiating $-BP$ once more contributes an additional factor $-B$. This is precisely the curvature term that Itô's formula detects.

Inserting all three derivatives into (4.6) gives

$$-(A' + B'r)P = -b_Q(r)BP + \tfrac{1}{2}v(r)B^2P - rP.$$

A zero-coupon price is positive; dividing through by P and reversing signs yields the *affine-pricing master identity*

$$\boxed{A'(\tau) + B'(\tau)r = r + B(\tau)b_Q(r) - \frac{1}{2}B(\tau)^2 v(r).} \tag{4.15}$$

It is useful to read this equation literally. The left-hand side is affine in the current state r. The first term on the right-hand side is affine. The other two terms involve the unknown pricing drift $b_Q(r)$ and variance $v(r)$, with *different maturity loadings*: B and B^2. This difference turns out to be what restricts their possible state dependence.

Why the convexity term matters. If we incorrectly used ordinary calculus on the exponential-affine bond price, the term $-\frac{1}{2}B^2 v(r)$ would disappear from (4.15). We would then derive the wrong maturity loading and, later, the wrong Riccati equation. Itô's correction is not an optional adjustment: it is the state diffusion acting on the curvature of the bond price.

4.5 Why two maturities constrain the pricing dynamics

The following result is an important selection in our compatibility hierarchy.

Proposition 4.1 (Pricing rigidity under an affine state representation)**.** *Suppose* (4.15) *holds for every interior state r and at two positive maturities τ_1, τ_2 for which $B_i = B(\tau_i)$ are distinct and nonzero. Then the risk-neutral drift $b_Q(r)$ and short-rate variance $v(r)$ are affine functions of r on that state interval.*

Proof. At each maturity τ_i, rearrange (4.15) and divide by B_i:

$$b_Q(r) - \frac{1}{2}B_i v(r) = \frac{A'(\tau_i)}{B_i} + \frac{B'(\tau_i) - 1}{B_i} r. \tag{4.16}$$

Each right-hand side is an affine function of r. Subtract the identity for $i = 2$ from that for $i = 1$:

$$\frac{B_2 - B_1}{2} v(r) = \left[\frac{A'(\tau_1)}{B_1} - \frac{A'(\tau_2)}{B_2}\right] + \left[\frac{B'(\tau_1) - 1}{B_1} - \frac{B'(\tau_2) - 1}{B_2}\right] r.$$

Because $B_2 - B_1 \neq 0$, dividing by this nonzero constant shows that $v(r)$ is affine. Substituting that conclusion back into either instance of (4.16) shows that $b_Q(r)$ is affine as well. □

The logic is worth emphasising. At one maturity, a nonlinear drift could in principle be offset by a nonlinear variance in (4.15). At *two* maturities with different B values, the weights of those two nonlinear functions differ. Exact cancellation at both maturities is possible only if both functions are themselves affine. This is a mathematical restriction, not an empirical claim that observed rates necessarily possess affine pricing dynamics.

Write the general affine coefficients as

$$b_Q(r) = m + \alpha r, \qquad v(r) = v_0 + \beta r. \tag{4.17}$$

Coefficient matching in (4.15) now gives

$$A' = mB - \frac{1}{2}v_0 B^2, \tag{4.18}$$

$$B' = 1 + \alpha B - \frac{1}{2}\beta B^2. \tag{4.19}$$

The B equation is a *Riccati equation*: a first-order ordinary differential equation whose right-hand side is quadratic in the unknown function. Notice that a variance intercept v_0 would change the A equation but not the B equation.

4.6 What selects square-root diffusion?

We now distinguish two very different statements: *the short rate is positive* and *the short-rate variance vanishes at the zero boundary.* Positivity of a state process alone does not imply that its local variance has zero intercept; state-domain choices and boundary behaviour matter. Our particular classification adds the condition

$$\lim_{r \downarrow 0} v(r) = v(0) = 0, \tag{4.20}$$

with continuous affine variance and $v(r) > 0$ for $r > 0$. Equation (4.17) then forces $v_0 = 0$ and $\beta > 0$. Thus

$$\boxed{v(r) = \beta r, \qquad \sigma(r) = \sqrt{\beta r}.} \tag{4.21}$$

A sufficient conventional condition for the square-root process with $m > 0$ to remain strictly positive when started in the positive interior is $2m \geq \beta$ (the familiar Feller boundary condition). Equality requires care in boundary classification; the important point is that the drift and boundary assumptions must ensure an admissible nonnegative solution, rather than assuming positivity from the square root alone.

Selection, not circular reasoning. Exact state-affine pricing plus two nondegenerate maturity loadings forces *affine* variance. The further requirement $v(0) = 0$ selects *square-root* variance. Neither Chapter 3's forecast consistency nor positivity by itself selects $v(r) = \beta r$.

Under these conditions, the pricing equations become

$$\boxed{A'(\tau) = mB(\tau), \qquad B'(\tau) = 1 + \alpha B(\tau) - \tfrac{1}{2}\beta B(\tau)^2, \qquad A(0) = B(0) = 0.} \tag{4.22}$$

This is our general one-state square-root *risk-neutral* affine pricing system. The sign of α has not yet been selected.

4.7 A careful look at the short-maturity boundary

The boundary values are helpful checks on all the signs. From $B(0) = 0$, the Riccati equation gives $B'(0) = 1$, as required by the short-rate boundary. Differentiating the Riccati equation gives

$$B''(\tau) = [\alpha - \beta B(\tau)]B'(\tau), \qquad B''(0) = \alpha. \tag{4.23}$$

Since $A' = mB$, we have $A''(0) = m$. The instantaneous forward rate is

$$x(\tau, r) = -\partial_\tau \log P(\tau, r) = mB(\tau) + B'(\tau)r. \tag{4.24}$$

Therefore

$$x(0, r) = r, \qquad \partial_\tau x(0, r) = m + \alpha r = b_Q(r). \tag{4.25}$$

The second equality is an especially useful identity. In this time-homogeneous one-factor model, the initial maturity slope of the forward curve equals the *risk-neutral* instantaneous short-rate drift. It is not, in general, the physical drift $b_P(r)$, nor an implication that the forward curve forecasts the realised short rate under $\mathbb{P}$.

For completeness, Taylor expansion near $\tau = 0$ gives

$$B(\tau) = \tau + \tfrac{1}{2}\alpha\tau^2 + O(\tau^3), \tag{4.26}$$

$$A(\tau) = \tfrac{1}{2}m\tau^2 + O(\tau^3), \tag{4.27}$$

$$x(\tau, r) = r + (m + \alpha r)\tau + O(\tau^2). \tag{4.28}$$

These expansions explain why the instantaneous-forward curve begins at the current short rate but acquires a maturity slope governed by pricing dynamics.

4.8 The Riccati equation is not yet a volatility-hump selection

It is tempting to move immediately to the positive-slope pricing branch $\alpha > 0$, because that is the branch ultimately used in our humped-volatility analysis. We deliberately do not do so in this chapter. The Riccati equation was obtained before asking for any hump.

Some qualitative conclusions can nevertheless be stated now. With $\beta > 0$ and an admissible nonexplosive positive branch, the positive equilibrium root of the Riccati equation is

$$B_\infty = \frac{\alpha + \sqrt{\alpha^2 + 2\beta}}{\beta} > 0. \tag{4.29}$$

Starting from $B(0) = 0$, the solution increases towards B_∞ and satisfies $B'(\tau) > 0$ at finite maturities. The maturity loading of instantaneous-forward volatility under $\mathbb{Q}$ is, directly from (4.24) and (4.21),

$$q(\tau, r) = \sigma(r)\partial_r x(\tau, r) = \sqrt{\beta r}\, B'(\tau). \tag{4.30}$$

Equation (4.23) shows that the sign of B'' is the sign of $\alpha - \beta B$. If $\alpha \leq 0$, the loading B' cannot have a strict interior maximum. If $\alpha > 0$, B' initially rises and later falls, producing a single interior hump. We shall prove that classification and derive its location in Chapter 5; no additional result about the *physical* short-rate drift follows from it.

4.9 Why rolling maturity must not be confused with a held bond

The earlier chapters distinguished $P(t, T)$, a particular bond maturing at fixed calendar date T, from $P_t(\tau) = P(t, t + \tau)$, the price quoted at a constant remaining maturity. We can make the distinction explicit. For a smooth price map $P(\tau, r_t)$, Itô's formula at fixed *rolling* τ gives

$$\mathrm{d}P_t(\tau) = \mathcal{L}_Q P(\tau, r_t)\,\mathrm{d}t + \sigma(r_t)\partial_r P(\tau, r_t)\,\mathrm{d}W_t^Q. \tag{4.31}$$

The pricing PDE implies $\mathcal{L}_Q P = \partial_\tau P + rP$, so the rolling-price drift is

$$\frac{\text{drift of } \mathrm{d}P_t(\tau)}{P_t(\tau)} = r_t + \frac{\partial_\tau P}{P} = r_t - x(\tau, r_t). \tag{4.32}$$

That drift is *not* required to equal r_t: maintaining constant maturity entails moving from one bond to another as time passes. The apparent discrepancy is simply the missing maturity roll-down.

For a fixed-date bond, remaining maturity changes according to $\mathrm{d}\tau = -\,\mathrm{d}t$. The chain rule adds $-\partial_\tau P\,\mathrm{d}t = +xP\,\mathrm{d}t$ to (4.31). The fixed-T drift therefore becomes exactly $r_t P$, as it must under $\mathbb{Q}$. More explicitly, using $\partial_r P = -BP$,

$$\boxed{\frac{\mathrm{d}P(t,T)}{P(t,T)} = r_t\,\mathrm{d}t - B(T-t)\sigma(r_t)\,\mathrm{d}W_t^Q.} \tag{4.33}$$

The Brownian coefficient is the same one that we shall later use in analysing bond expected excess returns. The maturity-transport correction, however, is essential for interpreting the drift.

Why this matters for compatibility. Exponentiating a forward rate and applying Itô's formula is mathematically valid, but does not by itself produce the value process of a traded bond. Exponentiating the *integrated* forward curve produces the bond price. Its fixed-calendar-maturity evolution is the object to which the traded-asset no-arbitrage drift condition applies.

4.10 A second route to the same pricing restrictions

There are two complementary approaches to arbitrage-free term-structure modelling. In the *short-rate pricing* approach used here, one starts with the risk-neutral generator and solves (4.6).

In the *HJM forward-rate* approach, one starts with forward-rate volatility and derives the drift required by no arbitrage. In the Musiela rolling-maturity coordinates, the one-Brownian HJM equation takes the form

$$\mathrm{d}x_t(\tau) = \left[\partial_\tau x_t(\tau) + q_t(\tau)\int_0^\tau q_t(u)\,\mathrm{d}u\right]\mathrm{d}t + q_t(\tau)\,\mathrm{d}W_t^Q. \tag{4.34}$$

A full derivation, beginning from the integrated forward curve and including Itô's convexity correction, will accompany the detailed HJM compatibility discussion. For present purposes, the relevant check is direct. In our affine model,

$$\begin{aligned}\text{state Itô drift of } x(\tau, r_t) &= B'(\tau)[m + \alpha r_t],\\ \partial_\tau x(\tau, r_t) &= mB'(\tau) + B''(\tau)r_t,\\ q(\tau, r_t) &= \sigma(r_t)B'(\tau),\\ \int_0^\tau q(u, r_t)\,\mathrm{d}u &= \sigma(r_t)B(\tau).\end{aligned}$$

Therefore the HJM drift identity reads

$$B'(m + \alpha r) = mB' + B''r + v(r)B'B. \tag{4.35}$$

For $v(r) = \beta r$, cancellation of the constant terms leaves

$$B'' = (\alpha - \beta B)B',$$

which is exactly the derivative of the Riccati equation (4.22). Hence the pricing-PDE and HJM routes agree. That agreement is a *compatibility check*, not a separate selection of the physical dynamics.

4.11 What Chapter 4 has established—and what remains free

We started from an affine curve representation selected by an explicit forecasting requirement, but not yet identified as arbitrage-free. Under an equivalent risk-neutral pricing measure and standard regularity conditions, the fixed-maturity traded-bond martingale restriction gave the pricing PDE. Its exponential-affine substitution yielded the master identity (4.15). Two distinct nonzero maturity loadings then forced the pricing drift and short-rate variance to be affine in the state.

Only after adding a zero-variance boundary did we obtain square-root volatility and the Riccati equations. We also derived the short-end identity $\partial_\tau x(0, r) = b_Q(r)$ and showed that rolling-maturity bond prices do not obey the same instantaneous drift rule as bonds with fixed maturity dates. The HJM and pricing-generator formulations agree under the resulting model.

What is still free? We have not selected $\alpha > 0$ (the hump branch); we have not selected the physical drift b_P, any physical stationary distribution, or the market price of risk. Nor have we asserted that any observed term structure is exactly one-state affine. Those are separate issues in the compatibility hierarchy.

The hierarchy after Chapter 4. Forecast closure (with its own regularity assumptions) gives a state-affine curve. Arbitrage-free pricing, plus two nondegenerate bond loadings, gives affine risk-neutral drift and variance. A separate zero-variance boundary gives square-root diffusion. The Riccati equations then fix the maturity loadings. Physical equilibrium and risk compensation have not yet been identified.

Bridge to Chapter 5

The next chapter studies the forward-volatility loading $B'(\tau)$ generated by the Riccati equation. We shall ask when it is decreasing, when it has a strict interior hump, and why selecting the humped *risk-neutral* branch does not preclude a restoring *physical* short-rate process.

Chapter 5

The Instantaneous-Forward Volatility Hump

Chapter 4 derived a family of exponential-affine bond prices and its Riccati equation without selecting the sign of the risk-neutral short-rate drift slope. This chapter asks a more specific question: under which conditions does the instantaneous-forward volatility loading *initially rise with maturity, reach one interior maximum, and subsequently decline? We derive the answer from the Riccati equation, explain the role of its two fixed points and obtain an exact hump location and peak-to-short-end ratio. We then distinguish the non-restoring* risk-neutral *drift selected by the hump from possible restoring* physical *drifts. The chapter establishes compatibility and its limits; it does not claim that all market forward-volatility curves must be humped or that a physical stationary distribution has already been selected.*

5.1 The question inherited from Chapter 4

Chapter 3 introduced an affine state-to-curve representation. Chapter 4 imposed no-arbitrage pricing and, after the additional zero-variance boundary selection, obtained the risk-neutral model

$$\mathrm{d}r_t = (m + \alpha r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^Q, \qquad m > 0, \quad \beta > 0, \quad \alpha \in \mathbb{R}. \tag{5.1}$$

For a positive starting state we use the conservative boundary condition $2m \geq \beta$, ensuring zero is unattainable in the familiar square-root setting; the strict positivity of the model is an additional modelling requirement, not a consequence of affinity alone. The zero-coupon price is

$$P(\tau, r) = \exp[-A(\tau) - B(\tau)r], \qquad A(0) = B(0) = 0, \tag{5.2}$$

where

$$A'(\tau) = mB(\tau), \qquad B'(\tau) = 1 + \alpha B(\tau) - \frac{\beta}{2}B(\tau)^2. \tag{5.3}$$

These equations did not determine the sign of α. Choosing it now will be a *further model selection*, based on the maturity geometry of forward-rate volatility.

What do we mean by a volatility hump? A forward rate with maturity τ can react more or less strongly to an infinitesimal shock to the same underlying short-rate state. An interior hump means that this *shock sensitivity* is smaller at the shortest maturity than at some intermediate maturity, but eventually decreases again at longer maturities. This is different from the statement that interest rates themselves, yields, or historical finite-period changes form a hump-shaped curve.

5.2 From bond prices to instantaneous-forward volatility

The instantaneous forward rate is defined by differentiating the logarithm of the bond price:

$$x(\tau, r) = -\partial_\tau \log P(\tau, r). \tag{5.4}$$

Substituting equation (5.2) and differentiating gives

$$\boxed{x(\tau, r) = A'(\tau) + B'(\tau)r = mB(\tau) + B'(\tau)r.} \tag{5.5}$$

At the short boundary, $B(0) = 0$ and $B'(0) = 1$, so $x(0, r) = r$ as required. Now hold the *remaining maturity coordinate* τ fixed and apply Itô's formula to $x(\tau, r_t)$. Because x is affine in r, its second state derivative is zero:

$$\partial_r x = B'(\tau), \qquad \partial_{rr} x = 0.$$

Thus the stochastic term is

$$\boxed{q(\tau, r_t)\, \mathrm{d}W_t^Q = \sqrt{\beta r_t}\, B'(\tau)\, \mathrm{d}W_t^Q, \qquad q(\tau, r) = \sqrt{\beta r}\, B'(\tau).} \tag{5.6}$$

The *instantaneous-forward volatility* q is a diffusion coefficient. For a very short time increment Δt, conditional on the state $r_t = r$,

$$\mathrm{Var}^Q[x_{t+\Delta t}(\tau) - x_t(\tau) \mid r_t = r] = q(\tau, r)^2 \Delta t + o(\Delta t), \tag{5.7}$$

with the same leading conditional quadratic variation under an equivalent physical measure. For fixed $r > 0$, every maturity shares the factor $\sqrt{\beta r}$; only $B'(\tau)$ determines the maturity *shape* of this instantaneous loading.

Four different objects. The instantaneous-forward volatility $q(\tau, r)$ is not (i) a hump in the level curve $x(\tau, r)$, (ii) the volatility of the zero-coupon yield, (iii) the volatility of a bond's holding return, or (iv) the unconditional standard deviation of forward-rate changes over a long observation window. These quantities have different maturity loadings, and the distinction matters for empirical tests.

5.3 Why the Riccati equation determines the hump

We can determine whether $B'(\tau)$ rises or falls without first solving the Riccati equation explicitly. Differentiate equation (5.3) with respect to τ:

$$\begin{aligned} B''(\tau) &= \alpha B'(\tau) - \beta B(\tau) B'(\tau) \\ &= \boxed{(\alpha - \beta B(\tau)) B'(\tau).} \end{aligned} \tag{5.8}$$

We must first justify the sign of B'. Define

$$\gamma_\alpha = \sqrt{\alpha^2 + 2\beta}, \qquad B_+ = \frac{\alpha + \gamma_\alpha}{\beta} > 0, \qquad B_- = \frac{\alpha - \gamma_\alpha}{\beta} < 0. \tag{5.9}$$

Then

$$B' = -\frac{\beta}{2}(B - B_+)(B - B_-). \tag{5.10}$$

Since $B(0) = 0$ lies between the roots and $B'(0) = 1$, the solution increases through the interval $0 < B(\tau) < B_+$ and tends to B_+ only as $\tau \to \infty$. Hence

$$B'(\tau) > 0 \quad \text{for every finite } \tau \geq 0. \tag{5.11}$$

We can therefore read the sign of B'' directly from $\alpha - \beta B$.

5.3.1 The restoring risk-neutral branch

Suppose $\alpha = -\kappa_Q$ with $\kappa_Q > 0$. The risk-neutral drift is then $m - \kappa_Q r$: for sufficiently high r it pulls the state back downward. But now

$$B''(\tau) = -(\kappa_Q + \beta B(\tau))B'(\tau) < 0. \tag{5.12}$$

Thus $B'(\tau)$ decreases from its initial value $B'(0) = 1$. There is no interior maximum. The boundary case $\alpha = 0$ also has $B''(\tau) < 0$ for $\tau > 0$.

The first restriction. A restoring *risk-neutral* affine short-rate drift in the selected square-root model cannot generate a strict interior hump in instantaneous-forward volatility. This says nothing yet about whether the *physical* drift is restoring.

5.3.2 The non-restoring risk-neutral branch

Suppose instead that $\alpha > 0$. To match the notation of our compatibility paper, write $\alpha = a > 0$. Equation (5.8) becomes

$$B'' = (a - \beta B)B'. \tag{5.13}$$

At the short end, $B(0) = 0$, so

$$B''(0) = a > 0. \tag{5.14}$$

The volatility loading initially *increases*. But $B(\tau)$ rises continuously towards

$$B_\infty = \frac{a + \sqrt{a^2 + 2\beta}}{\beta} > \frac{a}{\beta}. \tag{5.15}$$

Hence B must cross a/β exactly once. Before the crossing, $B'' > 0$; after the crossing, $B'' < 0$. We have established both existence and uniqueness of the hump.

Proposition 5.1 (Unique interior hump)**.** *Under equations* (5.1)–(5.3)*, $m, \beta > 0$, and the positive-state square-root specification, a strict interior maximum of $q(\tau, r)$ as a function of τ at each fixed $r > 0$ exists if and only if $\alpha > 0$. In that case it occurs at the unique maturity $\tau_* > 0$ satisfying*

$$\boxed{B(\tau_*) = \frac{a}{\beta}.} \tag{5.16}$$

The volatility loading is increasing before τ_* and decreasing thereafter.

Proof. Equation (5.11) ensures that B' is strictly positive. The sign of its derivative B'' is therefore the sign of $\alpha - \beta B$. If $\alpha \leq 0$ it cannot be positive for any positive maturity. If $\alpha = a > 0$, B increases from zero past a/β towards the larger fixed point B_∞, crossing once. The sign of B'' changes once from positive to negative. Since $q = \sqrt{\beta r}B'$ with $r > 0$ fixed, q has the same maturity maximum. □

5.4 How large is the hump?

The proof identifies the location through B, but also gives the *height* without solving an equation. At the peak substitute $B = a/\beta$ into the Riccati equation:

$$\begin{aligned} B'(\tau_*) &= 1 + a\frac{a}{\beta} - \frac{\beta}{2}\left(\frac{a}{\beta}\right)^2 \\ &= \boxed{1 + \frac{a^2}{2\beta}.} \end{aligned} \tag{5.17}$$

At the short boundary $B'(0) = 1$. Therefore the peak-to-short-maturity instantaneous volatility ratio is

$$\boxed{R := \frac{q(\tau_*, r)}{q(0, r)} = 1 + \frac{a^2}{2\beta} > 1.} \tag{5.18}$$

The short-rate state r cancels. So does the intercept m. The ratio depends only on the competition between the positive risk-neutral drift slope a and the variance slope β.

The limiting cases are instructive. As $a \downarrow 0$ for fixed β, $R \downarrow 1$ and the hump disappears into the short boundary. A greater value of a^2/β produces a more pronounced peak. This ratio alone, however, does not identify a and β separately.

5.5 Finding the hump maturity step by step

The hump condition $B(\tau_*) = a/\beta$ also yields an explicit maturity. Set

$$\gamma = \sqrt{a^2 + 2\beta}, \quad B_+ = \frac{a + \gamma}{\beta}, \quad B_- = \frac{a - \gamma}{\beta}. \tag{5.19}$$

Separating the Riccati equation gives

$$\frac{\mathrm{d}B}{(B_+ - B)(B - B_-)} = \frac{\beta}{2}\,\mathrm{d}\tau. \tag{5.20}$$

Since $B_+ - B_- = 2\gamma/\beta$, integration from $B(0) = 0$ to $B(\tau)$ produces

$$\log \frac{(B(\tau) - B_-)/(B_+ - B(\tau))}{(-B_-)/B_+} = \gamma\tau. \tag{5.21}$$

Equivalently, the solution can be written

$$\boxed{B(\tau) = \frac{2(e^{\gamma\tau} - 1)}{(\gamma - a)(e^{\gamma\tau} - 1) + 2\gamma}.} \tag{5.22}$$

Substituting $B = a/\beta$ and simplifying gives

$$\boxed{\tau_* = \frac{1}{\gamma} \log\left(\frac{\gamma + a}{\gamma - a}\right).} \tag{5.23}$$

This expression is positive because $\gamma > a > 0$. It shows that the hump location carries information about the maturity scale γ^{-1} as well as the ratio a/γ. In contrast to the peak-height ratio, its dimensions are those of time.

For reference, one may differentiate (5.22) directly:

$$B'(\tau) = \frac{4\gamma^2 e^{\gamma\tau}}{\left[(\gamma - a)e^{\gamma\tau} + (\gamma + a)\right]^2}. \tag{5.24}$$

This starts at one, reaches the peak in (5.17), and tends to zero as $\tau \to \infty$. The large-maturity decline is not an assumption about how volatility should decay; it follows from the Riccati solution within the chosen model.

5.6 A transparent numerical illustration

To see the shape without tying it to any empirical calibration, take dimensionally consistent illustrative parameters $a = 0.20\,\mathrm{yr}^{-1}$ and $\beta = 0.08\,\mathrm{yr}^{-2}$ when the short rate is expressed in inverse-year units. Then

$$\gamma = \sqrt{0.20^2 + 2(0.08)} \simeq 0.4472\,\mathrm{yr}^{-1}, \qquad R = 1 + \frac{0.20^2}{2(0.08)} = 1.25.$$

Equation (5.23) gives

$$\tau_* \simeq \frac{1}{0.4472} \log\left(\frac{0.4472 + 0.20}{0.4472 - 0.20}\right) \simeq 2.15 \text{ years}.$$

Thus, at any fixed positive state, the instantaneous forward's diffusion loading starts at $\sqrt{\beta r}$, rises by 25 percent to a single maximum around 2.15 years, and then falls towards zero. This is an *illustration*, not an estimate of a market volatility surface. The parameter dimensions follow from measuring the short rate in inverse years and calendar time in years.

5.7 The risk-neutral drift is non-restoring—what does that mean?

The hump selects

$$\boxed{\mathrm{d}r_t = (m + ar_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^Q, \qquad a > 0.} \tag{5.25}$$

For $r > 0$, its conditional instantaneous expected change under $\mathbb{Q}$ is $m + ar > 0$. Indeed, the conditional mean satisfies the ordinary differential equation

$$\frac{\mathrm{d}}{\mathrm{d}h}\mathbb{E}_r^Q[r_h] = m + a\mathbb{E}_r^Q[r_h], \qquad \mathbb{E}_r^Q[r_0] = r. \tag{5.26}$$

Solving it gives

$$\boxed{\mathbb{E}_r^Q[r_h] = \left(r + \frac{m}{a}\right)e^{ah} - \frac{m}{a}.} \tag{5.27}$$

The mean grows as the risk-neutral forecasting horizon increases. In particular, there is no *restoring invariant probability law under* $\mathbb{Q}$ for this branch. This is not the same as saying bond prices fail to exist.

The bond price depends on the discounted path functional

$$P(\tau, r) = \mathbb{E}_r^Q\left[\exp\left(-\int_0^\tau r_s\,\mathrm{d}s\right)\right], \tag{5.28}$$

not on the undiscounted mean (5.27). The Riccati solution $B(\tau)$ is bounded, and $A'(\tau) = mB(\tau)$ implies

$$\lim_{\tau\to\infty} -\frac{\log P(\tau, r)}{\tau} = mB_\infty = m\frac{a + \gamma}{\beta}. \tag{5.29}$$

Thus the long *yield* can remain finite even though the ordinary risk-neutral conditional mean grows. This distinction between risk-neutral expected states and risk-neutral discounted values is central to correct interpretation.

A common mistaken inference. The selected non-restoring drift under $\mathbb{Q}$ does *not* mean that interest rates must drift upwards in actual historical data, or that the physical process cannot possess a stationary distribution. The probability measures answer different questions.

5.8 How physical mean reversion can coexist with the hump

Now return to the physical measure. The diffusion coefficient of the same state process is unchanged by an equivalent Girsanov change of measure, but its drift can differ. Write

$$\mathrm{d}r_t = b_P(r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^P. \tag{5.30}$$

Adopt the convention already used in Chapter 4 and the SSRN paper:

$$\boxed{b_Q(r) = b_P(r) + \sqrt{\beta r}\,\lambda(r).} \tag{5.31}$$

Once $b_Q(r) = m + ar$ is fixed by pricing, this becomes

$$\boxed{\lambda(r) = \frac{m + ar - b_P(r)}{\sqrt{\beta r}}.} \tag{5.32}$$

It is an identity between the two drifts and the market price of risk. It does not by itself prescribe b_P, and a formally defined λ is not automatically a valid equivalent change of measure: integrability and martingale conditions are required on the intended horizon.

Consider a simple admissible physical example:

$$b_P(r) = m - \kappa_P r, \qquad \kappa_P > 0. \tag{5.33}$$

Then

$$\lambda(r) = \frac{a + \kappa_P}{\sqrt{\beta}}\sqrt{r}. \tag{5.34}$$

Under $\mathbb{P}$ the drift restores the state towards m/κ_P, while under $\mathbb{Q}$ it increases with r. On finite horizons, under standard regularity, the market-price-of-risk adjustment can reconcile the two specifications. They have *the same instantaneous diffusion coefficient*. The physical CIR example is therefore compatible with the selected humped *pricing* branch; it is not selected uniquely by it.

There is an even broader possibility. For any suitable restoring physical drift b_P admitting an invariant distribution and a valid measure change, equation (5.32) defines the risk-price function required by the chosen pricing drift. The function b_P need not be affine, and the physical stationary density need not be Gamma. This freedom is one of the principal conclusions of our framework, but the full conditions on stationary densities and boundary behaviour will be developed in later chapters.

The coexistence mechanism. Humped forward volatility constrains the *pricing* drift slope: $a > 0$. Physical stationarity constrains the *physical* drift and its invariant distribution. The risk-price function λ connects the two, allowing physical restoring dynamics and non-restoring risk-neutral dynamics to describe the same diffusion under different measures.

5.9 What the hump does—and does not—identify

The exact ratio $R = 1 + a^2/(2\beta)$ and hump location τ_* provide information about the pricing parameters. The intercept m does not enter either expression; it affects the level of the forward curve and long yield, but not the normalised maturity profile of B'. In principle, knowledge of R and τ_* under the exact model can help identify a and β. In practice, instantaneous-forward volatility is inferred rather than directly observed, and estimation errors can be significant.

More importantly, even perfect knowledge of $B'(\tau)$ and of the pricing parameters does *not* identify $b_P(r)$ or $\lambda(r)$ separately. We know their difference through (5.31), but not the two individual functions. Conditional physical forecasts, expected bond excess returns or an independently estimated physical stationary density would supply further information. That is where the later chapters of the monograph will extend the compatibility hierarchy.

Scope of the result. The if-and-only-if sign selection of Proposition 5.1 holds *within* the one-state square-root affine model with the specified boundary and time-homogeneous Riccati pricing equation. It is not a universal theorem that all humped forward-volatility curves require non-restoring risk-neutral rates. Multifactor models, alternative state domains and time-dependent specifications can exhibit different geometries.

5.10 What Chapter 5 has established—and what remains free

We began with the Riccati pricing equations of Chapter 4 and found the forward volatility loading $q(\tau, r) = \sqrt{\beta r}\, B'(\tau)$. The identity $B'' = (\alpha - \beta B)B'$ then showed that a restoring risk-neutral affine drift gives a decreasing loading, whereas a strictly positive risk-neutral drift slope selects a *unique interior hump*. For that selected branch, we derived its location, its peak ratio and the long-maturity decay of the volatility loading. None of these calculations assumed physical stationarity.

We then separated the measures. A non-restoring $\mathbb{Q}$ drift is compatible with restoring $\mathbb{P}$ dynamics, including but not limited to physical CIR, provided the proposed measure change and diffusion are admissible. *What remains free?* The physical drift, physical stationary law, and market price of risk have not been uniquely determined. Nor have we imposed monotonicity of the physical equilibrium forward curve or a long-forward mean restriction.

The hierarchy after Chapter 5. Forecast closure selected a state-affine curve; no-arbitrage and the positive zero-variance boundary selected square-root affine pricing; the additional strict hump requirement now selects the non-restoring *risk-neutral* drift branch. It does not collapse the family of compatible *physical* drifts to CIR.

Bridge to Chapter 6

Chapter 6 will examine the distinction between physical and pricing measures in greater detail. We shall derive the role of the market price of risk, explain the change-of-measure conditions, and show why physical forecasts and risk-neutral valuations can differ while remaining compatible. This will prepare the derivation of expected excess bond returns and, eventually, physical stationary equilibrium.

Chapter 6

Physical and Risk-Neutral Measures

Chapter 5 showed that a hump in instantaneous-forward volatility selects a non-restoring risk-neutral short-rate drift within the chosen one-state affine square-root model. This need not imply increasing short rates under the physical probability law. This chapter explains why: arbitrage-free prices and physical forecasts are calculated under distinct probability measures, linked by the market price of risk. We derive the drift transformation with a consistent sign convention, distinguish instantaneous drift differences from finite-horizon forecast differences, and demonstrate how physically restoring dynamics can coexist with a non-restoring pricing drift. We also establish an important stationary-mean identity and explain its relation to the short-end slope of the physical equilibrium forward curve. A detailed treatment of zero-coupon bond excess returns follows in Chapter 7.

6.1 Why two probability measures are necessary

The short rate is an economic variable that evolves through calendar time. Under the *physical* measure $\mathbb{P}$, probabilities describe the frequencies or beliefs used for forecasting. Under a *risk-neutral* pricing measure $\mathbb{Q}$, probabilities are adjusted so that correctly discounted traded asset prices have no predictable excess return, subject to the standard absence-of-arbitrage and numeraire assumptions.

The two measures need not assign the same probabilities to future short-rate outcomes. This does *not* introduce two separate short rates: there is one state process r_t, viewed under two probability laws. In particular, the market price of risk is not a second independently observed interest-rate process.

Recall the model selected in Chapter 5:

$$\mathrm{d}r_t = (m + ar_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^Q, \qquad m > 0,\ a > 0,\ \beta > 0. \tag{6.1}$$

The coefficient $a > 0$ is needed for the interior hump in the *instantaneous-forward volatility loading*. The physical drift is not determined by this selection. Write, at interior states $r > 0$,

$$\mathrm{d}r_t = b_P(r_t)\,\mathrm{d}t + \sigma(r_t)\,\mathrm{d}W_t^P, \qquad \sigma(r) = \sqrt{\beta r}. \tag{6.2}$$

We now derive the relation between b_P and the risk-neutral drift $b_Q(r) = m + ar$.

6.2 A small-time picture of a change of measure

For a short interval h, conditioned on $r_t = r$, the diffusion has the leading-order physical behaviour

$$r_{t+h} - r = b_P(r)h + \sigma(r)(W_{t+h}^P - W_t^P) + o_P(\sqrt{h}). \tag{6.3}$$

The conditional random fluctuation has scale $\sigma(r)\sqrt{h}$; the conditional average change has scale h. Changing the probabilities used to value risk changes the drift—the conditional average rate of change—without changing the instantaneous quadratic variation of the same continuous process.

This is a restricted claim: it concerns *equivalent changes of measure* for the same diffusion on a specified time horizon. It does not say that every alternative stochastic model has the same volatility, nor that equivalent measures imply the same finite-horizon variances.

The interpretation. The market price of risk tells us how the pricing probability law tilts the likelihood of stochastic shocks. The short-rate drift changes because those shocks are assigned different probabilities, even though the underlying state and its local shock amplitude remain the same.

6.3 Girsanov's theorem with an explicit sign convention

Let W^P be a Brownian motion under $\mathbb{P}$, and let $\lambda_t = \lambda(r_t)$ be a progressively measurable process. Define

$$Z_t = \exp\left(\int_0^t \lambda_s \, \mathrm{d}W_s^P - \frac{1}{2}\int_0^t \lambda_s^2 \, \mathrm{d}s\right). \tag{6.4}$$

Assume Z_t is a true martingale on the chosen finite horizon $[0, T]$, with $\mathbb{E}^P[Z_T] = 1$, and define

$$\left.\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right|_{\mathcal{F}_T} = Z_T. \tag{6.5}$$

One familiar *sufficient*, but not necessary, condition is Novikov's condition

$$\mathbb{E}^P\left[\exp\left(\frac{1}{2}\int_0^T \lambda_s^2 \, \mathrm{d}s\right)\right] < \infty. \tag{6.6}$$

Under these conditions Girsanov's theorem states that

$$\boxed{W_t^Q = W_t^P - \int_0^t \lambda_s \, \mathrm{d}s} \tag{6.7}$$

is Brownian motion under $\mathbb{Q}$. Differentiating this notation gives

$$\mathrm{d}W_t^P = \mathrm{d}W_t^Q + \lambda_t \, \mathrm{d}t. \tag{6.8}$$

Substitute equation (6.8) into the physical SDE (6.2):

$$\begin{aligned} \mathrm{d}r_t &= b_P(r_t)\,\mathrm{d}t + \sigma(r_t)(\,\mathrm{d}W_t^Q + \lambda(r_t)\,\mathrm{d}t) && (6.9)\\ &= \big[b_P(r_t) + \sigma(r_t)\lambda(r_t)\big]\,\mathrm{d}t + \sigma(r_t)\,\mathrm{d}W_t^Q. && (6.10) \end{aligned}$$

Comparing with the risk-neutral SDE yields the identity

$$\boxed{b_Q(r) = b_P(r) + \sigma(r)\lambda(r).} \tag{6.11}$$

This sign convention is maintained throughout the monograph. Alternative texts sometimes define the market price of risk with the opposite sign; what matters is consistent use of the density, the Brownian-motion shift, and the drift relation.

A technical limit to the argument. Writing down $\lambda = (b_Q - b_P)/\sigma$ pointwise does not by itself prove that the resulting measures are equivalent. The stochastic exponential must define a probability measure on the horizon in question, and the SDEs must meet the appropriate existence and boundary requirements. In the positive square-root setting, $\sigma(0) = 0$ makes boundary behaviour especially important. Localisation and non-explosion can establish finite-horizon equivalence under conditions weaker than Novikov's condition. No equivalence on the infinite-horizon sigma-field is asserted here.

6.4 Why the generator changes but the local variance does not

Chapter 2 introduced the physical generator

$$\mathcal{L}_P\phi(r) = b_P(r)\phi'(r) + \frac{1}{2}\sigma^2(r)\phi''(r). \tag{6.12}$$

Under $\mathbb{Q}$, the same twice differentiable test function has generator

$$\mathcal{L}_Q\phi(r) = b_Q(r)\phi'(r) + \frac{1}{2}\sigma^2(r)\phi''(r). \tag{6.13}$$

Subtracting these operators and using equation (6.11) gives

$$\boxed{(\mathcal{L}_Q - \mathcal{L}_P)\phi(r) = \sigma(r)\lambda(r)\phi'(r).} \tag{6.14}$$

Notice what cancels: the second-derivative, or Itô convexity, term. Equivalent measure changes modify the drift of functions of the state through the *first* derivative. They do not remove or alter the instantaneous quadratic variation term.

For $\phi(r) = r$, we have $\phi' = 1$ and $\phi'' = 0$, so

$$(\mathcal{L}_Q - \mathcal{L}_P)r = \sigma(r)\lambda(r). \tag{6.15}$$

For $\phi(r) = r^2$, the difference is instead

$$(\mathcal{L}_Q - \mathcal{L}_P)r^2 = 2r\sigma(r)\lambda(r). \tag{6.16}$$

These examples reinforce a point we will use later: a change of probability measure is an operator relationship, not merely a single numerical adjustment to one expected value.

6.5 Local forecast differences versus finite-horizon forecasts

Define, for a common initial state $r_t = r$,

$$m_P(h, r) = \mathbb{E}_r^P[r_{t+h}], \qquad m_Q(h, r) = \mathbb{E}_r^Q[r_{t+h}]. \tag{6.17}$$

For small h, the generator identity implies, subject to the same regularity and integrability conditions used in Chapter 2,

$$m_P(h, r) = r + b_P(r)h + o(h), \tag{6.18}$$

$$m_Q(h, r) = r + b_Q(r)h + o(h). \tag{6.19}$$

Subtracting,

$$\boxed{\partial_h[m_Q(h, r) - m_P(h, r)]|_{h=0} = \sigma(r)\lambda(r).} \tag{6.20}$$

This is the precise meaning of the statement that risk compensation shifts the risk-neutral short-rate *forecast drift*. It is not correct to add the rate $\sigma\lambda$ directly to a future short-rate *level* without integrating its effects in time.

The distinction persists at finite horizons. Define the physical and pricing semigroups

$$(\mathsf{P}_h\phi)(r) = \mathbb{E}_r^P[\phi(r_h)], \qquad (\mathsf{Q}_h\phi)(r) = \mathbb{E}_r^Q[\phi(r_h)]. \tag{6.21}$$

They satisfy their respective backward evolution equations with generators $\mathcal{L}_P$ and $\mathcal{L}_Q$. Even if their generators differ by the simple multiplication–differentiation operator in equation (6.14), their finite-horizon forecasts ordinarily have different evolving state distributions.

There is an informative exact representation for the *linear* Q drift in our selected model. Under either law, define $b_Q(r) = m + ar$ as a function of state. Applying Itô's formula to $e^{-as}r_s$, and taking physical conditional expectations, yields

$$\boxed{\mathbb{E}_r^P[r_h] = e^{ah}r + \frac{m}{a}(e^{ah} - 1) - \int_0^h e^{a(h-s)}\mathbb{E}_r^P[\sigma(r_s)\lambda(r_s)]\,\mathrm{d}s.} \tag{6.22}$$

Under $\mathbb{Q}$ the corresponding risk compensation term is absent:

$$\boxed{\mathbb{E}_r^Q[r_h] = e^{ah}r + \frac{m}{a}(e^{ah} - 1).} \tag{6.23}$$

Consequently,

$$\boxed{\mathbb{E}_r^Q[r_h] - \mathbb{E}_r^P[r_h] = \int_0^h e^{a(h-s)}\mathbb{E}_r^P[\sigma(r_s)\lambda(r_s)]\,\mathrm{d}s.} \tag{6.24}$$

The integral formula uses the affine Q drift and appropriate moment conditions; it is not a generic identity for arbitrary nonlinear b_Q. Importantly, the expectation in the integrand is under $\mathbb{P}$ and is taken along the *physical* state distribution. This is a convenient finite-horizon extension of the local relationship, not the claim that the two measures are interchangeable.

6.6 A worked example: restoring P and non-restoring Q

Consider the *illustrative* physical specification

$$\mathrm{d}r_t = (m - \kappa_P r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^P, \qquad \kappa_P > 0. \tag{6.25}$$

Chapter 5 selected the risk-neutral drift $b_Q(r) = m + ar$, with $a > 0$. The drift difference is

$$b_Q(r) - b_P(r) = (a + \kappa_P)r, \tag{6.26}$$

and hence

$$\boxed{\lambda(r) = \frac{a + \kappa_P}{\sqrt{\beta}}\sqrt{r}.} \tag{6.27}$$

This premium is regular as $r \downarrow 0$ and is unbounded as $r \to \infty$; the latter behaviour is not inherently a mathematical problem, but does mean that global equivalence must be justified rather than inferred from boundedness. Both square-root diffusions are nonexplosive on finite horizons, and under the strict-positive-boundary condition $2m \geq \beta$ their laws, started from the same positive state, can be compared by localisation.

Now compute the two conditional means directly:

$$\mathbb{E}_r^P[r_h] = \frac{m}{\kappa_P} + \left(r - \frac{m}{\kappa_P}\right)e^{-\kappa_P h}, \tag{6.28}$$

$$\mathbb{E}_r^Q[r_h] = \left(r + \frac{m}{a}\right)e^{ah} - \frac{m}{a}. \tag{6.29}$$

Under $\mathbb{P}$, the conditional mean returns to m/κ_P. Under $\mathbb{Q}$, the conditional mean rises with the horizon. They are different descriptions of the same short-rate state under different probability weights.

Why there is no paradox. A stationary physical short-rate process can support a humped instantaneous-forward volatility loading and a non-restoring risk-neutral drift. A state-dependent market price of risk changes the pricing probabilities enough to reverse the sign of the state-drift slope. This construction establishes *possibility*, not empirical plausibility or uniqueness of the physical drift.

The physical CIR example has an invariant Gamma distribution, with mean m/κ_P, under standard recurrence and boundary conditions. But there is no reason to require (6.25) for the general compatibility analysis: suitable nonlinear restoring physical drifts can also coexist with the fixed affine pricing drift. The stationary-density classification is developed later in the monograph.

6.7 Physical equilibrium and the mean premium identity

The distinction between the measures becomes especially transparent at physical equilibrium. Suppose the $\mathbb{P}$ diffusion admits an invariant probability distribution $\pi_P(r)\,\mathrm{d}r$ with finite first moment

$$\theta_P = \mathbb{E}_{\pi_P}[r] < \infty, \tag{6.30}$$

and assume sufficient integrability and generator-domain conditions to apply stationarity to the test function $\phi(r) = r$. Since the law is invariant, its expected short rate does not change with calendar time:

$$0 = \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}_{\pi_P}[r_t] = \mathbb{E}_{\pi_P}[b_P(r)]. \tag{6.31}$$

Take the *physical stationary expectation* of the *pointwise* drift identity (6.11):

$$\mathbb{E}_{\pi_P}[b_Q(r)] = \mathbb{E}_{\pi_P}[b_P(r)] + \mathbb{E}_{\pi_P}[\sigma(r)\lambda(r)]. \tag{6.32}$$

By stationarity, the first term on the right vanishes. With $b_Q(r) = m + ar$,

$$\boxed{\mathbb{E}_{\pi_P}[\sigma(r)\lambda(r)] = m + a\theta_P > 0.} \tag{6.33}$$

The mean is positive on the humped branch ($m, a, \theta_P > 0$); this does *not* assert that $\lambda(r)$ is positive at every state.

Which expectation is this? The quantity $\mathbb{E}_{\pi_P}[b_Q(r)]$ means that the function b_Q has been averaged over the *physical stationary density*. It is *not* the same object as $\mathbb{E}^Q[b_Q(r_t)]$, and it does not assert that $\mathbb{Q}$ possesses a stationary distribution. The distinction is particularly important here because the selected Q branch is non-restoring.

6.8 Why the equilibrium forward curve initially rises

Our affine forward curve from Chapter 5 is

$$x(\tau, r) = mB(\tau) + B'(\tau)r, \tag{6.34}$$

where $B(0) = 0$, $B'(0) = 1$, and $B''(0) = a$ from the Riccati equation. Average this curve over the physical stationary law:

$$x_e^P(\tau) = \mathbb{E}_{\pi_P}[x(\tau, r)] = mB(\tau) + \theta_P B'(\tau). \tag{6.35}$$

Differentiating with respect to maturity gives

$$\partial_\tau x_e^P(\tau) = mB'(\tau) + \theta_P B''(\tau). \tag{6.36}$$

At $\tau = 0$,

$$\boxed{\partial_\tau x_e^P(0) = m + a\theta_P.} \tag{6.37}$$

Combining equations (6.33) and (6.37), we obtain

$$\boxed{\partial_\tau x_e^P(0) = \mathbb{E}_{\pi_P}[b_Q(r)] = \mathbb{E}_{\pi_P}[\sigma(r)\lambda(r)].} \tag{6.38}$$

This compact identity is one of the clearest expressions of compatibility in the present model. Physical stationarity removes the average *physical* drift; it does not make the stationary expected forward curve flat. Its initial maturity slope is governed by the average pricing-drift adjustment.

A useful interpretation. The non-restoring risk-neutral drift is not a prediction that real-world interest rates continually rise. At physical equilibrium, its positive stationary average is exactly the mean risk-compensation adjustment that separates physical forecasting from pricing. In Chapter 7 we will identify this same quantity with the stationary mean *instantaneous expected excess holding-return rate* of a zero-coupon bond, after multiplication by its Riccati loading $B(\tau)$.

6.9 What this chapter establishes—and what remains free

The mathematical chain is now explicit. The hump selects $b_Q(r) = m + ar$ within the model; the change of measure provides

$$\boxed{b_P(r) = m + ar - \sqrt{\beta r}\,\lambda(r).} \tag{6.39}$$

Physical equilibrium can coexist with this pricing structure whenever the selected physical drift admits an admissible invariant law and the finite-horizon change of measure is valid. Under the requisite moment conditions, stationarity gives the mean identity (6.33), which also equals the initial maturity slope of the stationary expected forward curve.

We have *not* deduced a unique $b_P(r)$, $\lambda(r)$, or stationary density from the hump and pricing model alone. We have not asserted $\mathbb{E}^Q[r_h] = \mathbb{E}^P[r_h] + \sigma\lambda$, nor assumed positivity of the market price of risk at every state. We have also not established that an empirically observed hump identifies this model uniquely.

The next chapter adds an especially useful piece of information: the conditional expected *excess holding return* of an actual zero-coupon bond. It will show how this expected return provides a maturity-dependent observation of the same drift difference $b_Q - b_P$, thereby sharpening the identification question without imposing a physical CIR law.

Chapter 7

Expected Excess Bond Returns and Physical Forecasts

The preceding chapters established affine bond pricing, a possible hump in instantaneous-forward volatility, and the change of measure connecting physical and risk-neutral short-rate dynamics. We now apply these results to the actual holding return on a zero-coupon bond. The key identity shows that the physical short-rate forecast drift and the instantaneous expected excess bond-return rate jointly reproduce the risk-neutral drift. We derive the relation from the bond-price stochastic differential equation, examine its state and maturity restrictions, extend it away from physical stationarity, and connect it to stationary forward-curve geometry. The resulting identities are standard consequences of affine pricing and measure change; their role here is to make the hierarchy of compatibility explicit, not to posit new pricing laws.

7.1 Why bond returns are the missing connection

There are two complementary descriptions of a given short-rate process. Under the physical probability measure $\mathbb{P}$, we ask how the short rate is expected to evolve. Under the risk-neutral measure $\mathbb{Q}$, we ask what prices of traded bonds are consistent with the absence of arbitrage. Chapter 6 established that these descriptions are connected by the market price of risk:

$$b_Q(r) = b_P(r) + \sigma(r)\lambda(r). \tag{7.1}$$

This equation alone does not identify b_P or λ: knowledge of their sum is not knowledge of each summand. But the expected holding return on a bond provides an economic interpretation of the difference. The same λ that changes the probability measure also determines compensation for bearing a bond's interest-rate exposure.

Our one-state risk-neutral model is

$$\mathrm{d}r_t = (m + ar_t)\,\mathrm{d}t + \sqrt{\beta r_t}\;\mathrm{d}W_t^Q, \qquad m, \beta > 0, \quad a > 0, \tag{7.2}$$

where $a > 0$ is the humped branch obtained in Chapter 5. We write $\sigma(r) = \sqrt{\beta r}$ and $b_Q(r) = m + ar$. To distinguish the general identity from the particular humped branch, some formulas below retain b_Q explicitly.

Under the pricing representation derived in Chapter 4, a zero-coupon bond with calendar maturity T has price

$$P(t,T) = \exp[-A(\tau) - B(\tau)r_t], \qquad \tau = T - t, \tag{7.3}$$

with

$$A'(\tau) = mB(\tau), \qquad B'(\tau) = 1 + aB(\tau) - \tfrac{1}{2}\beta B(\tau)^2, \qquad A(0) = B(0) = 0. \tag{7.4}$$

For $\tau > 0$, $B(\tau) > 0$. This positive loading will matter repeatedly.

Terminology. Throughout the chapter, $\mathrm{ERP}_\tau(r)$ means the *physical conditional instantaneous expected excess holding-return rate* on a zero-coupon bond with *remaining maturity* τ. It is not the bond's yield spread, a conventional term premium, or the realised return over a month or a year. We must derive this rate using the P&L of a bond held with its calendar maturity T fixed.

7.2 The return on a bond held through calendar time

This is where the rolling-versus-fixed-maturity distinction from Chapter 1 becomes indispensable. The function in (7.3) is conveniently written in the rolling coordinate $\tau = T - t$; the investor, however, holds a claim with a fixed maturity *date* T. Thus, along the bond's life,

$$\mathrm{d}\tau = -\,\mathrm{d}t. \tag{7.5}$$

Let $F(\tau, r) = e^{-A(\tau)-B(\tau)r}$. Its elementary derivatives are

$$F_r = -BF, \qquad F_{rr} = B^2F, \tag{7.6}$$

$$F_\tau = -(A' + B'r)F. \tag{7.7}$$

Applying Itô's formula to $F(T - t, r_t)$ *under the physical measure* gives

$$\mathrm{d}P(t,T) = -F_\tau\,\mathrm{d}t + F_r\,\mathrm{d}r_t + \tfrac{1}{2}F_{rr}\sigma^2(r_t)\,\mathrm{d}t \tag{7.8}$$

$$= P(t,T)\Big[A' + B'r_t - Bb_P(r_t) + \tfrac{1}{2}B^2\sigma^2(r_t)\Big]\,\mathrm{d}t - P(t,T)B\sigma(r_t)\,\mathrm{d}W_t^P. \tag{7.9}$$

Here A, B, B' are all evaluated at $\tau = T - t$. Therefore

$$\frac{\mathrm{d}P(t,T)}{P(t,T)} = \Big[A' + B'r_t - Bb_P(r_t) + \tfrac{1}{2}B^2\beta r_t\Big]\,\mathrm{d}t - B\sqrt{\beta r_t}\,\mathrm{d}W_t^P. \tag{7.10}$$

The affine pricing master equation from Chapter 4 states

$$A' + B'r = r + Bb_Q(r) - \tfrac{1}{2}B^2\beta r. \tag{7.11}$$

Substituting this into (7.10) yields a substantial simplification:

$$\boxed{\frac{\mathrm{d}P(t,T)}{P(t,T)} = \big[r_t + B(\tau)(b_Q(r_t) - b_P(r_t))\big]\,\mathrm{d}t - B(\tau)\sigma(r_t)\,\mathrm{d}W_t^P.} \tag{7.12}$$

The stochastic term has conditional mean zero, under the usual integrability assumptions. The predictable instantaneous excess of the bond's holding-return rate over the short rate is consequently

$$\boxed{\mathrm{ERP}_\tau(r) = B(\tau)[b_Q(r) - b_P(r)] = B(\tau)\sigma(r)\lambda(r).} \tag{7.13}$$

In the adopted sign convention, the bond's Brownian exposure is $-B\sigma$ while its expected excess return is $+B\sigma\lambda$. This sign is consistent with the choice $\mathrm{d}W^Q = \mathrm{d}W^P - \lambda\,\mathrm{d}t$ in Chapter 6; using an alternative market-price-of-risk convention would change both signs consistently.

Economic interpretation. The bond's price sensitivity $B(\tau)$ tells us how much the bond reacts to the short-rate shock. The quantity $\sigma(r)\lambda(r)$ specifies the risk-price adjustment to the short-rate drift. Their product gives the instantaneous excess holding-return rate. It is not necessary to treat an exponential of a single forward rate as a traded security: the traded bond price already provides the appropriate P&L object.

7.3 The pointwise forecast–return compatibility identity

We can now state the central result in a particularly transparent form.

Proposition 7.1 (Pointwise forecast–return compatibility)**.** *Suppose the affine bond pricing equations* (7.3)–(7.4) *and the finite-horizon change of measure in Chapter 6 hold. For each interior short-rate state* $r > 0$ *and maturity* $\tau > 0$*,*

$$\boxed{\mathrm{ERP}_\tau(r) + B(\tau)b_P(r) = B(\tau)b_Q(r) = B(\tau)(m + ar).} \tag{7.14}$$

No physical stationarity is required.

Proof. From (7.13), $\mathrm{ERP}_\tau = B(b_Q - b_P)$. Adding Bb_P to both sides gives (7.14). The equation holds state by state, not merely after averaging over a stationary distribution. □

It is useful to read (7.14) in both directions. Given the physical drift and the risk-neutral pricing model, it determines the instantaneous expected bond excess return. Conversely, if the state dependence of that expected return were known independently, then

$$\boxed{b_P(r) = m + ar - \frac{\mathrm{ERP}_\tau(r)}{B(\tau)}.} \tag{7.15}$$

In this second reading, the same physical drift must be recovered from *every* positive maturity. This is a model-identification statement conditional on the ability to estimate instantaneous expected returns; it is not a claim that the returns are directly observable.

7.3.1 Why curvature must cancel

In the humped affine branch $b_Q(r) = m + ar$. Assuming b_P and ERP_τ are twice differentiable in the state, differentiate (7.14) twice while holding maturity fixed:

$$\boxed{\frac{\partial^2}{\partial r^2}\,\mathrm{ERP}_\tau(r) = -B(\tau)b_P''(r).} \tag{7.16}$$

Thus any nonlinearity in the physical forecast drift is offset by opposite state curvature in expected excess returns. In particular, an affine expected-return function at one positive maturity forces b_P to be affine. The implication is *conditional on that further restriction on returns*: affine pricing by itself permits nonlinear physical drifts.

7.3.2 A cross-maturity restriction

For any two maturities $\tau_1, \tau_2 > 0$, equation (7.13) implies

$$\boxed{\frac{\mathrm{ERP}_{\tau_1}(r)}{B(\tau_1)} = \frac{\mathrm{ERP}_{\tau_2}(r)}{B(\tau_2)} = m + ar - b_P(r).} \tag{7.17}$$

This is stronger than requiring expected excess returns to have some similar shape: their entire state dependence must be identical after division by the appropriate maturity loading. The restriction applies *pointwise in* r. It is a standard one-factor consequence of affine bond-return pricing, here made explicit as a cross-maturity compatibility condition. At zero maturity $B(0) = 0$, so the ratios in (7.17) are intentionally stated only for positive maturities.

7.4 Linking bond returns to the forward-curve spread

We can obtain a second representation of the same compatibility relationship from the forward curve. Recall that

$$x(\tau, r) = mB(\tau) + B'(\tau)r, \qquad x(0, r) = r. \tag{7.18}$$

Subtract the short rate, then apply the Riccati equation:

$$\begin{aligned} x(\tau, r) - r &= mB + (B' - 1)r && (7.19)\\ &= mB + \left(aB - \tfrac{1}{2}\beta B^2\right) r && (7.20)\\ &= B\big[m + ar\big] - \tfrac{1}{2}\beta r B^2. && (7.21) \end{aligned}$$

Using $m + ar = b_P(r) + \sigma(r)\lambda(r)$ and (7.13) finally gives

$$\boxed{x(\tau, r) - r = \underbrace{\mathrm{ERP}_\tau(r)}_{\text{risk compensation}} + \underbrace{B(\tau) b_P(r)}_{\text{physical drift contribution}} - \underbrace{\tfrac{1}{2}\beta r B(\tau)^2}_{\text{pricing-volatility correction}}}. \tag{7.22}$$

This identity is exact within the model. The left-hand side is an *instantaneous forward-rate spread* over the current short rate, not the yield spread on a bond; likewise the first term on the right is an instantaneous expected holding-return premium, not a conventional yield term premium. Equation (7.22) shows how risk compensation, physical forecast drift, and convexity jointly generate the observed forward curve.

7.5 Averaging away from equilibrium

The pointwise results above do not depend on stationarity. Let $\mu_t = \mathbb{E}^P[r_t]$ denote the mean short rate under an arbitrary physical law at calendar time t. Assume the relevant first moments and generator-domain conditions ensure

$$\dot{\mu}_t = \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}^P[r_t] = \mathbb{E}^P[b_P(r_t)]. \tag{7.23}$$

Take physical expectations on both sides of (7.22), keeping maturity τ fixed when defining the curve being averaged. Because $B(\tau)$ is deterministic under our time-homogeneous affine model,

$$\boxed{\mathbb{E}^P[x(\tau, r_t)] - \mu_t = \mathbb{E}^P[\mathrm{ERP}_\tau(r_t)] + B(\tau)\dot{\mu}_t - \tfrac{1}{2}\beta\mu_t B(\tau)^2.} \tag{7.24}$$

The middle term describes the *instantaneous* drift of the physical mean. It is not equal to the entire physical forecast over a finite horizon. Chapter 6 showed that finite-horizon forecast differences involve an integral of risk compensation along the evolving physical distribution; equation (7.24) instead concerns a cross-sectional forward curve at one calendar time.

What changes outside equilibrium? The physical mean need not be constant, so the forward-curve spread contains a nonzero $B\dot{\mu}_t$ contribution. The pricing geometry itself remains the same as long as the assumed time-homogeneous risk-neutral model remains valid. There is no need to invoke physical equilibrium to obtain the return–forecast identity.

7.6 The stationary equilibrium relation

Suppose now that r_t has an invariant physical distribution π_P with finite first moment $\theta_P = \mathbb{E}_{\pi_P}[r]$ and that the expected drift vanishes:

$$\mathbb{E}_{\pi_P}[b_P(r)] = 0. \tag{7.25}$$

The stationary physical mean forward curve is

$$x_e^P(\tau) = \mathbb{E}_{\pi_P}[x(\tau, r)] = mB(\tau) + \theta_P B'(\tau). \tag{7.26}$$

Taking the stationary expectation of (7.13) gives, by (7.1) and (7.25),

$$\boxed{\mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau] = B(\tau)\mathbb{E}_{\pi_P}[\sigma(r)\lambda(r)] = B(\tau)(m + a\theta_P).} \tag{7.27}$$

Taking the stationary expectation of (7.22) gives the promised equilibrium forward-spread decomposition:

$$\boxed{x_e^P(\tau) - \theta_P = \mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau] - \tfrac{1}{2}\beta\theta_P B(\tau)^2.} \tag{7.28}$$

This formula exhibits the expected-return contribution and pricing-volatility correction separately. Neither term is generally itself a quoted yield term premium.

There is a particularly transparent short-maturity boundary identity. Differentiating (7.26), using $B(0) = 0$, $B'(0) = 1$ and $B''(0) = a$, yields

$$\boxed{\partial_\tau x_e^P(\tau)\Big|_{\tau=0} = m + a\theta_P = \mathbb{E}_{\pi_P}[\sigma(r)\lambda(r)].} \tag{7.29}$$

Because $a, m, \theta_P > 0$ on the selected branch, this stationary expected risk-price adjustment is positive. That does *not* require $\lambda(r)$, or instantaneous expected excess returns, to be positive at every state. Combining (7.27) and (7.29) gives

$$\boxed{\mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau] = B(\tau)\,\partial_\tau x_e^P(0).} \tag{7.30}$$

The entire maturity dependence of stationary *average instantaneous* bond excess returns is fixed by the Riccati loading and the short-end slope of the stationary mean forward curve.

7.7 What would identify the physical stationary distribution?

Equation (7.15) shows that an independently identified state-dependent ERP function would reveal b_P. Under further conditions ensuring a positive recurrent diffusion and zero stationary probability current, the stationary Fokker–Planck equation then provides the invariant density.

For $\sigma^2(r) = \beta r$, the stationary zero-current condition is

$$b_P(r)\pi_P(r) - \tfrac{1}{2}\frac{\mathrm{d}}{\mathrm{d}r}[\beta r \pi_P(r)] = 0. \tag{7.31}$$

Dividing by $\pi_P(r) > 0$ and rearranging gives

$$b_P(r) = \tfrac{\beta}{2}\left[1 + r\frac{\mathrm{d}}{\mathrm{d}r}\log\pi_P(r)\right]. \tag{7.32}$$

If $H(r) = \mathrm{ERP}_\tau(r)/B(\tau)$ (the same function for every positive maturity), then $b_P(r) = m + ar - H(r)$, so

$$\boxed{\pi_P(r) = C\, r^{2m/\beta - 1} \exp\left[\frac{2a}{\beta}r - \frac{2}{\beta}\int^r \frac{H(u)}{u}\,\mathrm{d}u\right].} \tag{7.33}$$

The lower integration reference point is arbitrary and absorbed into the normalising constant C. Formula (7.33) is a *candidate* density: its integral must be finite, its boundary behaviour admissible, and the associated diffusion positive recurrent. Formal substitution alone does not demonstrate that these conditions hold.

7.7.1 Physical CIR as a worked example, not a requirement

Suppose we choose the additional restoring affine drift

$$b_P(r) = m - \kappa_P r, \qquad \kappa_P > 0. \tag{7.34}$$

Then

$$H(r) = (a + \kappa_P)r, \qquad \mathrm{ERP}_\tau(r) = (a + \kappa_P)B(\tau)r. \tag{7.35}$$

Substitution into (7.33) yields

$$\pi_P(r) = C r^{2m/\beta - 1} e^{-2\kappa_P r/\beta}, \tag{7.36}$$

which is Gamma with shape $2m/\beta$ and rate $2\kappa_P/\beta$. This example makes the inverse map explicit. But $H(r)$ need not be proportional to r, and in general the invariant density need not be Gamma. Physical stationarity does not itself impose the affine choice (7.34).

7.7.2 Stationary expected-return distributions across maturities

Under a common stationary short-rate law, the random conditional expected-return rate satisfies

$$\mathrm{ERP}_\tau(r) = B(\tau)H(r), \qquad r \sim \pi_P. \tag{7.37}$$

Thus the distributions at distinct maturities form a common *scale family*. Wherever the moments exist, their means scale as $B(\tau)$, their variances as $B(\tau)^2$, and their standardised skewness and kurtosis are the same. This is a property of *conditional expected return rates regarded as random functions of a stationary state*, not necessarily of realised finite-period bond returns. It does not select a Gamma, lognormal, or other particular family by itself.

7.8 The empirical and logical limits

The identities in this chapter distinguish compatibility from identification. First, $\mathrm{ERP}_\tau(r)$ is a conditional instantaneous expectation, not a directly observable realised return. Its empirical estimation requires data, a forecast model, and choices concerning time variation and measurement error. Second, a realised return over a fixed holding horizon is affected by the evolution of the short rate and the decline in the bond's remaining maturity; it is not obtained merely by multiplying the instantaneous ERP at the initial state by the holding period. Third, the identification formula (7.15) relies on the specified risk-neutral affine model and its known $B(\tau)$; uncertainty in those parameters transmits into the inferred physical drift.

Finally, one should not infer that the present section has discovered a novel no-arbitrage formula. The bond factor-loading times the market price of risk is a standard affine term-structure identity; see, for example, Duffee (2002) and Piazzesi's survey on affine term-structure models. Our contribution is the explicit organisation of its consequences within the forecast–pricing–equilibrium hierarchy.

7.9 What we have established—and what remains free

Established under the selected model. The holding return of a fixed-calendar-maturity bond determines the instantaneous excess-return rate $B\sigma\lambda$. Its sum with Bb_P is Bb_Q at each state, and its state dependence is identical across positive maturities after scaling by B. The

forward-spread identity holds without physical stationarity; at physical equilibrium it yields the stationary expected-return and short-boundary slope relations.

Not established without further information. These equations do not determine the physical drift from pricing alone, select CIR or Gamma equilibrium, enforce pointwise positive risk compensation, or identify conventional yield term premia. Reconstructing a stationary density requires additional information about the conditional return function and the diffusion's boundary behaviour.

The compatibility lesson. Restrictive risk-neutral pricing geometry and flexible physical forecasting dynamics are reconciled by risk compensation. Neither physical drift nor expected excess returns can vary independently of the other; the model fixes their combined contribution. This is how established pricing identities become illuminating when organised into a single compatibility framework.

Chapter 8

Physical Equilibrium and Stationary Distributions

The preceding chapters have constrained risk-neutral bond pricing and connected physical short-rate drift to expected excess bond returns. This chapter explains what it means for the short rate to possess a physical stationary distribution. Beginning with the generator and conservation of probability, we derive the forward Kolmogorov equation, probability current, and zero-current invariant density. We make the necessary distinction between a formal stationary solution and an admissible invariant probability law, show how the stationary density determines the physical drift, and construct both CIR/Gamma and genuinely non-CIR examples compatible with the same humped affine pricing geometry. The chapter closes by connecting the stationary-density score to the compatibility coefficients and by identifying what remains undetermined.

8.1 Why equilibrium is a probability distribution, not a constant rate

Chapter 7 established, for a positive remaining maturity τ,

$$\mathrm{ERP}_\tau(r) = B(\tau)\big[b_Q(r) - b_P(r)\big], \qquad b_Q(r) = m + ar, \quad \sigma^2(r) = \beta r. \tag{8.1}$$

Knowing the pricing drift and factor loading does not, by itself, determine the physical drift b_P. But once we have specified a physical drift we can ask a different question: can the resulting short rate fluctuate indefinitely while maintaining the same probability distribution?

A stationary distribution π_P is an invariant probability law: if r_0 has distribution π_P, then r_t has the same distribution for every $t \geq 0$. Individual rates continue to move. In particular, equilibrium does *not* require $\mathrm{d}r_t = 0$ path by path or a constant forward curve. It refers to the law of the stochastic state, not to the elimination of randomness.

We work on the state space $(0, \infty)$ with the physical diffusion

$$\mathrm{d}r_t = b_P(r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^P, \qquad \beta > 0. \tag{8.2}$$

For the formal calculations, assume sufficient smoothness and integrability. When constructing an invariant law, we additionally require a well-defined, conservative diffusion with appropriate boundary behaviour. These requirements are not implied by writing down a nonnegative function that solves a differential equation.

8.2 From the generator to the forward Kolmogorov equation

In Chapter 2, Itô's formula gave the physical infinitesimal generator

$$\mathcal{L}_P\phi(r) = b_P(r)\phi'(r) + \frac{\beta r}{2}\phi''(r). \tag{8.3}$$

It describes how the *expected value of a test function* initially changes when the process starts at a given state. If $p(t, r)$ denotes the probability density of r_t, then

$$\mathbb{E}[\phi(r_t)] = \int_0^\infty \phi(r)p(t, r)\,\mathrm{d}r. \tag{8.4}$$

Under conditions permitting differentiation beneath the integral and use of the generator,

$$\frac{\mathrm{d}}{\mathrm{d}t}\int_0^\infty \phi(r)p(t, r)\,\mathrm{d}r = \int_0^\infty \Big[b_P(r)\phi'(r) + \frac{\beta r}{2}\phi''(r)\Big]p(t, r)\,\mathrm{d}r. \tag{8.5}$$

To transfer the derivatives from ϕ to p, first integrate by parts:

$$\int_0^\infty b_P\phi' p\,\mathrm{d}r = [\phi b_P p]_0^\infty - \int_0^\infty \phi\,\partial_r(b_P p)\,\mathrm{d}r. \tag{8.6}$$

The second-derivative term requires two integrations by parts:

$$\frac{1}{2}\int_0^\infty \beta r\phi'' p\,\mathrm{d}r = \frac{1}{2}[\beta r\phi' p]_0^\infty - \frac{1}{2}\int_0^\infty \phi'\,\partial_r(\beta r p)\,\mathrm{d}r \tag{8.7}$$

$$= \frac{1}{2}[\beta r\phi' p - \phi\,\partial_r(\beta r p)]_0^\infty + \frac{1}{2}\int_0^\infty \phi\,\partial_{rr}(\beta r p)\,\mathrm{d}r. \tag{8.8}$$

For test functions with compact support inside $(0, \infty)$, these boundary contributions vanish automatically. Equation (8.5) therefore implies the interior *forward Kolmogorov*, or *Fokker–Planck*, equation

$$\boxed{\partial_t p(t, r) = -\partial_r\big(b_P(r)p(t, r)\big) + \frac{1}{2}\partial_{rr}\big(\beta r\, p(t, r)\big).} \tag{8.9}$$

The backward generator $\mathcal{L}_P$ acts on functions of the state; its forward adjoint $\mathcal{L}_P^*$ acts on probability densities. They describe the same stochastic system from complementary perspectives.

8.3 Probability current and conservation of probability

The Fokker–Planck equation can be rewritten as a conservation equation:

$$\partial_t p(t, r) = -\partial_r J(t, r), \qquad J(t, r) = b_P(r)p(t, r) - \frac{1}{2}\partial_r\big(\beta r\, p(t, r)\big). \tag{8.10}$$

Here J is a *probability current*: positive J represents net transport of probability toward higher short rates. For an interval (u, v),

$$\frac{\mathrm{d}}{\mathrm{d}t}\Pr(u < r_t < v) = J(t, u) - J(t, v). \tag{8.11}$$

This formula gives the conservation interpretation precisely: probability can leave an interval only by crossing its boundaries.

A stationary density $\pi_P(r)$ satisfies $\partial_t p = 0$, hence

$$\partial_r J_\pi(r) = 0, \qquad J_\pi(r) = J_0 \tag{8.12}$$

for some constant J_0. A *zero-current* stationary law has $J_0 = 0$. This is the relevant closure for the conservative, non-circulating positive short-rate diffusions considered here, with no probability injected at one boundary and removed at the other. Stationarity alone, as an interior differential equation, says only that the current is constant; the zero value uses the boundary conditions.

Boundary conditions matter. One cannot infer $J_0 = 0$ merely by setting $\partial_t p = 0$ in the interior. For a diffusion on an interval with a source, a sink, or a reinjection mechanism, a stationary nonzero current may occur. Our invariant-density formula below is conditional on the conservative zero-flux boundary regime.

8.4 Deriving the zero-current stationary density

Set $J_\pi = 0$ in (8.10):

$$b_P(r)\pi_P(r) = \frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}r}\big(\beta r \pi_P(r)\big). \tag{8.13}$$

Expanding the derivative gives

$$b_P(r)\pi_P(r) = \frac{\beta}{2}\big[\pi_P(r) + r\pi_P'(r)\big]. \tag{8.14}$$

At states where $\pi_P(r) > 0$, divide by $\pi_P(r)$:

$$\boxed{b_P(r) = \frac{\beta}{2}\left[1 + r\,\partial_r \log \pi_P(r)\right].} \tag{8.15}$$

This reverses the usual modelling direction. Instead of first choosing a physical drift and then solving for its stationary distribution, we may begin with a candidate equilibrium density and calculate the physical drift compatible with it, subject to diffusion admissibility.

Alternatively, isolate the derivative of $\log \pi_P$:

$$\frac{\pi_P'(r)}{\pi_P(r)} = \frac{2b_P(r)}{\beta r} - \frac{1}{r}. \tag{8.16}$$

Integrate from an arbitrary reference state $r_* > 0$ to r:

$$\log \frac{\pi_P(r)}{\pi_P(r_*)} = \int_{r_*}^{r} \frac{2b_P(u)}{\beta u}\,\mathrm{d}u - \log \frac{r}{r_*}. \tag{8.17}$$

Exponentiating, and absorbing factors depending on r_* into a constant, yields

$$\boxed{\pi_P(r) = \frac{C}{\beta r}\exp\left[\int_{r_*}^{r} \frac{2b_P(u)}{\beta u}\,\mathrm{d}u\right], \qquad r > 0.} \tag{8.18}$$

One may absorb the constant β into C. We display it to emphasise the more general diffusion formula $\pi \propto \sigma^{-2}\exp\int 2b/\sigma^2$.

Proposition 8.1 (Zero-current drift–density correspondence)**.** *Suppose the physical diffusion* (8.2) *admits a positive, differentiable stationary density π_P with zero probability current and the integrations above are justified. Then* (8.15) *and* (8.18) *hold. Conversely, if a prescribed drift gives a positive, integrable solution of* (8.18) *and the associated diffusion is conservative with boundary conditions supporting zero flux, the normalised solution is an invariant density.*

The converse includes the diffusion and boundary hypotheses deliberately. A normalisable expression is not, on its own, a proof of positive recurrence or even the existence of the required nonexplosive process.

8.5 Formal solutions, recurrence and the zero boundary

There are three different questions which should not be conflated.

Normalisation. Does the candidate stationary density satisfy

$$0 < \int_0^\infty \frac{1}{\beta r} \exp\left[\int_{r_*}^r \frac{2b_P(u)}{\beta u}\,\mathrm{d}u\right]\,\mathrm{d}r < \infty? \tag{8.19}$$

If yes, a unique normalising constant C can be chosen.

Boundary admissibility. Does the SDE define a conservative process on the chosen state space, with a boundary condition consistent with the intended equilibrium? A density behaving well for $r > 0$ does not settle what happens if r_t reaches 0. Accessible reflecting, absorbing, entrance and inaccessible boundaries are mathematically different.

Long-run behaviour. Is the diffusion positive recurrent, so that its invariant probability law genuinely describes persistent long-run equilibrium? For a regular irreducible one-dimensional diffusion, finite speed measure together with the relevant nonexplosion and boundary conditions gives the familiar route to a stationary probability law. A full boundary classification uses the scale and speed measures; we give their essential forms below.

For a general diffusion $\mathrm{d}r = b(r)\,\mathrm{d}t + \sigma(r)\,\mathrm{d}W$, the scale density and speed density, up to irrelevant positive constants, are

$$s'(r) = \exp\left[-\int_{r_*}^r \frac{2b(u)}{\sigma^2(u)}\,\mathrm{d}u\right], \qquad m_{\text{speed}}(r) = \frac{2}{\sigma^2(r)s'(r)}. \tag{8.20}$$

The zero-current invariant density is proportional to the speed density when its total mass is finite and the boundary regime permits an invariant probability. These expressions clarify the connection between normalisation, recurrence and boundaries, without attempting a full classification theorem here.

In the CIR-type specification $\sigma^2(r) = \beta r$ with $b_P(0) = m > 0$, the local comparison near zero is with the classical square-root diffusion. The familiar sufficient condition

$$2m \geq \beta \tag{8.21}$$

ensures that, starting strictly positive, the CIR zero boundary is not reached. When $0 < 2m < \beta$, zero can be accessible without being absorbing if the inward drift and model boundary condition are retained. Nonnegativity, strict positivity, boundary inaccessibility and positive recurrence are distinct claims. Further conditions at large r are needed to prevent explosion and establish stationarity.

8.6 Worked example I: CIR and the Gamma density

Take restoring affine physical drift

$$b_P(r) = m - \kappa_P r, \qquad \kappa_P > 0. \tag{8.22}$$

Insert this into (8.18):

$$\pi_P(r) = \frac{C}{\beta r} \exp\left[\int^r \left(\frac{2m}{\beta u} - \frac{2\kappa_P}{\beta}\right)\,\mathrm{d}u\right] \tag{8.23}$$

$$= C' r^{2m/\beta - 1} e^{-2\kappa_P r/\beta}. \tag{8.24}$$

After normalisation,

$$\boxed{r \sim \text{Gamma}\left(\frac{2m}{\beta}, \frac{2\kappa_P}{\beta}\right),} \tag{8.25}$$

where the second Gamma parameter is a *rate*, not a scale. Its mean and variance are

$$\theta_P = \mathbb{E}_{\pi_P}[r] = \frac{m}{\kappa_P}, \qquad \text{Var}_{\pi_P}(r) = \frac{m\beta}{2\kappa_P^2}. \tag{8.26}$$

Gamma stationarity is therefore a consequence of *this particular* restoring affine physical drift and the square-root variance, not of arbitrage-free affine pricing by itself. The risk-neutral drift may simultaneously be $m + ar$, with $a > 0$, provided a suitable admissible change of measure connects the two laws.

8.7 Worked example II: a non-CIR physical equilibrium

To see that physical equilibrium need not be Gamma, retain the same square-root variance and pricing drift but consider

$$\boxed{b_P(r) = m - \kappa_P r - \delta\frac{r}{1+r}, \qquad \kappa_P > 0, \quad \delta > 0.} \tag{8.27}$$

The extra restoring term is nonlinear in r. Near the origin it vanishes, leaving $b_P(0) = m$ and the same local square-root boundary structure as in the CIR example. At large r, the linear negative term dominates, providing a restoring force and controlling the upper tail.

Inserting (8.27) into the zero-current equation, the logarithmic derivative is

$$\frac{\mathrm{d}}{\mathrm{d}r}\log\pi_P(r) = \left(\frac{2m}{\beta} - 1\right)\frac{1}{r} - \frac{2\kappa_P}{\beta} - \frac{2\delta}{\beta(1+r)}. \tag{8.28}$$

Each term integrates elementarily:

$$\boxed{\pi_P(r) = C\, r^{2m/\beta - 1} e^{-2\kappa_P r/\beta}(1+r)^{-2\delta/\beta}, \quad r > 0.} \tag{8.29}$$

The density is integrable near zero because $2m/\beta > 0$, and at infinity because of the decaying exponential. It is *not* a Gamma density for $\delta > 0$. Under the usual nonexplosion and inward/zero-flux boundary conditions (in particular, a strictly positive boundary is ensured when $2m \geq \beta$), this gives an admissible non-CIR physical stationary model.

Yet the risk-neutral pricing structure is unchanged:

$$\sigma(r)\lambda(r) = b_Q(r) - b_P(r) = (a + \kappa_P)r + \delta\frac{r}{1+r}. \tag{8.30}$$

Consequently, the expected instantaneous excess return is

$$\text{ERP}_\tau(r) = B(\tau)\left[(a+\kappa_P)r + \delta\frac{r}{1+r}\right]. \tag{8.31}$$

This is nonlinear in the state, but it has exactly the same maturity loading $B(\tau)$ as the CIR example. It is an explicit instance of the curvature cancellation discussed in Chapter 7. Establishing an equivalent change of measure for a chosen time horizon requires the usual martingale or equivalent-diffusion conditions; the algebraic drift identity by itself is not a substitute for those conditions.

A concrete compatibility example. The same humped affine bond-pricing model can coexist with a Gamma stationary physical law or with the non-Gamma law (8.29). The difference is absorbed in the state dependence of risk compensation, not in the risk-neutral Riccati equation.

8.8 Stationary moments and the meaning of equilibrium

We can now connect the density description to conditional forecasting. For any sufficiently regular test function ϕ whose generator is integrable under π_P and whose boundary terms vanish,

$$\boxed{\mathbb{E}_{\pi_P}[\mathcal{L}_P\phi(r)] = 0.} \tag{8.32}$$

This is the weak expression of stationarity. With $\phi(r) = r$, equation (8.3) gives $\mathcal{L}_P r = b_P(r)$, and therefore, assuming a finite first moment,

$$\boxed{\mathbb{E}_{\pi_P}[b_P(r)] = 0.} \tag{8.33}$$

This does *not* mean $b_P(r) = 0$ at every state. At low states the physical drift may be positive, and at high states negative; its stationary average vanishes.

For $\phi(r) = r^2$,

$$\mathcal{L}_P r^2 = 2rb_P(r) + \beta r. \tag{8.34}$$

If the second moment and associated integrability conditions hold, stationarity implies

$$\boxed{2\mathbb{E}_{\pi_P}[rb_P(r)] + \beta\theta_P = 0.} \tag{8.35}$$

The first and second moment identities constrain the physical law but generally do not specify the complete invariant density. In the non-CIR example, equation (8.33) becomes

$$m = \kappa_P\theta_P + \delta\mathbb{E}_{\pi_P}\left[\frac{r}{1+r}\right]. \tag{8.36}$$

There is no reason for the mean to equal m/κ_P when $\delta > 0$. That familiar expression belongs specifically to CIR.

At physical stationarity, Chapter 6 gave

$$\mathbb{E}_{\pi_P}[\sigma(r)\lambda(r)] = \mathbb{E}_{\pi_P}[b_Q(r) - b_P(r)] = m + a\theta_P. \tag{8.37}$$

Thus physical equilibrium and positive stationary average risk compensation can coexist with a non-restoring Q drift. The expectation on the left is taken under the *physical* stationary law; no Q-stationary distribution is being assumed.

8.9 The stationary-density score and the compatibility coefficients

Define the logarithmic-density score

$$S_P(r) = \partial_r \log \pi_P(r). \tag{8.38}$$

Using (8.15) and $\sigma(r) = \sqrt{\beta r}$,

$$\sigma'(r) = \frac{\beta}{2\sigma(r)}, \qquad \frac{b_P(r)}{\sigma(r)} = \frac{\beta}{2\sigma(r)} + \frac{\sigma(r)}{2}S_P(r). \tag{8.39}$$

Consequently the physical compatibility coefficient introduced in the SSRN framework may be written

$$\boxed{\mathcal{K}_X(r) := \sigma'(r) - \frac{b_P(r)}{\sigma(r)} = -\frac{\sigma(r)}{2}\partial_r \log \pi_P(r).} \tag{8.40}$$

Define the corresponding pricing coefficient $\mathcal{K}_x(r) := \sigma'(r) + \lambda(r)$. Subtracting the two yields

$$\boxed{\mathcal{K}_x(r) - \mathcal{K}_X(r) = \lambda(r) + \frac{b_P(r)}{\sigma(r)} = \frac{b_Q(r)}{\sigma(r)}.} \tag{8.41}$$

This identity again explains both the strength and the limitation of compatibility. The difference of coefficients is fixed by pricing, while their individual values depend on the physical density or market price of risk. For a known admissible $\mathcal{K}_X$, one may integrate (8.40) to construct the candidate density

$$\pi_P(r) = C \exp\left[-2 \int_{r_*}^{r} \frac{\mathcal{K}_X(u)}{\sigma(u)}\, \mathrm{d}u\right], \tag{8.42}$$

again subject to normalisation and the diffusion's boundary conditions.

8.10 What we have established—and what remains free

Established under the zero-current, admissible physical diffusion. The Fokker–Planck equation describes conservation of probability; the stationary density determines the physical drift through its logarithmic derivative; and the stationary generator identities imply vanishing average physical drift. We have exhibited both Gamma and non-Gamma physical stationary laws compatible at the level of drift and volatility with the same non-restoring risk-neutral affine pricing branch.

Not established by pricing or stationarity alone. Neither the exact physical drift nor the stationary density is uniquely selected. Nor does formal normalisation alone prove a valid conservative diffusion, guarantee an equivalent P–Q change of measure, or show that the risk-neutral process itself is stationary. These questions require separate boundary, recurrence and change-of-measure hypotheses.

The compatibility lesson. The pricing hierarchy constrains the bond-price geometry, but physical equilibrium is a property of the physical diffusion's invariant law. A family of admissible stationary distributions can coexist with a single pricing structure because different physical drifts are reconciled by correspondingly different risk-compensation functions. Chapter 9 will use this distinction to derive the physical equilibrium forward curve and examine the additional, explicitly optional, geometry selection.

Chapter 9

The Physical Equilibrium Forward Curve

Chapter 8 constructed admissible physical stationary laws without selecting a unique physical drift. We now ask what those laws imply for the average forward-rate curve. Exact state affinity means that a stationary expected forward curve depends on the physical invariant law through its mean alone. We derive its maturity slope from the Riccati equation, connect its spread to stationary expected instantaneous bond excess returns, and distinguish physical stationarity from an optional requirement that the mean forward curve increase at every maturity. The latter selection imposes a sharp bound on the physical mean. We derive the bound directly and describe its long-forward-measure interpretation, anticipating Chapter 10. Worked examples illustrate why stationarity, a forward-volatility hump and a nondecreasing equilibrium curve are distinct conditions.

9.1 From a stationary short rate to an expected forward curve

Chapter 8 showed that an admissible physical diffusion may possess an invariant density π_P even when its risk-neutral pricing drift is non-restoring. Let

$$\mathrm{d}r_t = b_P(r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^P, \qquad \theta_P := \int_0^\infty r\pi_P(r)\,\mathrm{d}r < \infty. \tag{9.1}$$

The notation θ_P denotes the *physical stationary mean*, not a parameter imposed on every possible physical drift. Stationarity means that starting with $r_0 \sim \pi_P$ leaves $r_t \sim \pi_P$ at all later dates. It does not imply that the rate is fixed at θ_P.

The previous pricing chapters selected the exponential-affine bond-price representation and its associated instantaneous-forward curve,

$$P(\tau, r) = \exp[-A(\tau) - B(\tau)r], \tag{9.2}$$

$$x(\tau, r) = mB(\tau) + B'(\tau)r. \tag{9.3}$$

Here $m, \beta > 0$ and, for the humped branch analysed in Chapter 5, $a > 0$. The loading satisfies

$$B'(\tau) = 1 + aB(\tau) - \frac{\beta}{2}B(\tau)^2, \quad B(0) = 0, \quad B'(0) = 1. \tag{9.4}$$

The physical measure determines the distribution of r_t; the $\mathbb{Q}$ pricing model determines the functions A and B. These roles must not be confused.

We define the *physical equilibrium forward curve* as the stationary average of the realised instantaneous-forward curves:

$$x_e^P(\tau) := \int_0^\infty x(\tau, r)\pi_P(r)\,\mathrm{d}r. \tag{9.5}$$

Substituting equation (9.3), with B and B' constant with respect to r, gives

$$x_e^P(\tau) = \int_0^\infty [mB(\tau) + B'(\tau)r]\pi_P(r)\,\mathrm{d}r \tag{9.6}$$

$$= mB(\tau)\underbrace{\int_0^\infty \pi_P(r)\,\mathrm{d}r}_{=1} + B'(\tau)\underbrace{\int_0^\infty r\pi_P(r)\,\mathrm{d}r}_{=\theta_P}. \tag{9.7}$$

Therefore

$$\boxed{x_e^P(\tau) = mB(\tau) + \theta_P B'(\tau) = x(\tau, \theta_P).} \tag{9.8}$$

The final equality is a consequence of *affinity*, not of stationarity by itself. In a nonlinear state-to-curve model, the mean forward curve generally depends on more than the stationary mean of the short rate.

A family of distinct stationary physical distributions can produce *exactly the same* equilibrium forward curve whenever they share the same finite mean θ_P and the same pricing functions. The average curve therefore cannot identify the full physical stationary density.

9.2 What equilibrium means at the short and long boundaries

At zero maturity, $B(0) = 0$ and $B'(0) = 1$, hence

$$\boxed{x_e^P(0) = \theta_P.} \tag{9.9}$$

The equilibrium forward curve begins at the stationary mean short rate. It does not generally begin at the current realised short rate: $x(0, r_t) = r_t$ holds for the realised curve, whereas equation (9.9) concerns an average over stationary states.

To understand the far end, define

$$\gamma := \sqrt{a^2 + 2\beta}, \qquad B_\infty := \frac{a+\gamma}{\beta}. \tag{9.10}$$

The Riccati solution satisfies $B(\tau) \uparrow B_\infty$ and $B'(\tau) \downarrow 0$ eventually as $\tau \to \infty$; on the humped branch B' first rises and then falls, but remains positive. It follows that

$$\boxed{\lim_{\tau\to\infty} x_e^P(\tau) = \rho := mB_\infty.} \tag{9.11}$$

The long-forward limit is independent of θ_P because the state loading $B'(\tau)$ vanishes. This says that the *long instantaneous forward* is deterministic within the model; it does not say that finite-maturity forwards or short rates are deterministic.

A convenient first description of the equilibrium geometry is consequently

$$x_e^P(0) = \theta_P, \qquad x_e^P(\infty) = \rho. \tag{9.12}$$

These endpoint identities do *not* by themselves imply that the curve is monotone between the endpoints. We must examine its derivative.

9.3 Deriving the maturity slope

Differentiate equation (9.8) with respect to remaining maturity:

$$\partial_\tau x_e^P(\tau) = mB'(\tau) + \theta_P B''(\tau). \tag{9.13}$$

The Riccati equation provides the second derivative without needing to solve B explicitly. Differentiate (9.4):

$$B''(\tau) = aB'(\tau) - \beta B(\tau)B'(\tau) \tag{9.14}$$

$$= \big[a - \beta B(\tau)\big]B'(\tau). \tag{9.15}$$

Inserting this into (9.13) gives

$$\boxed{\partial_\tau x_e^P(\tau) = B'(\tau)\big[m + a\theta_P - \beta\theta_P B(\tau)\big].} \tag{9.16}$$

Because $B'(\tau) > 0$ at every finite maturity, the sign of the equilibrium slope is determined entirely by the bracket.

At the short boundary $B(0) = 0$, so

$$\boxed{\partial_\tau x_e^P(0) = m + a\theta_P.} \tag{9.17}$$

For our humped branch, $m, a, \theta_P > 0$ and the initial slope is positive. Yet the bracket in equation (9.16) declines as $B(\tau)$ rises. A positive initial slope does not guarantee a nondecreasing equilibrium curve at every maturity.

> The hump of *instantaneous-forward volatility* and a turning point of the *mean forward curve* are different phenomena. The volatility maximum occurs where $B = a/\beta$; the mean-curve slope, if it becomes zero, does so where $B = (m + a\theta_P)/(\beta\theta_P)$. The latter threshold is strictly greater than a/β when $m > 0$.

9.4 Why physical stationarity does not select monotonicity

A physically stationary law requires $\mathbb{E}_{\pi_P}[b_P(r)] = 0$ when the identity function lies in the appropriate generator domain and its drift is integrable. This is a statement about the physical *time* evolution of the short rate. By contrast, $\partial_\tau x_e^P(\tau)$ measures the slope across *maturities* of an average pricing curve. There is no general reason these two derivatives should coincide.

The stationarity calculation in Chapter 8 permits many values of θ_P depending on the physical drift. Therefore equation (9.16) permits equilibrium curves that rise everywhere, or rise initially and eventually decline toward ρ. Both configurations can be compatible with a humped instantaneous-forward volatility loading.

If we wish to select only a globally nondecreasing stationary mean forward curve, we must state that as an *additional geometric assumption*:

$$\partial_\tau x_e^P(\tau) \geq 0 \qquad \text{for every } \tau \geq 0. \tag{9.18}$$

It is not an implication of the definition of physical equilibrium, of the CIR example, or of the no-arbitrage pricing condition.

Proposition 9.1 (Global monotonicity criterion)**.** *Assume $m, \beta > 0$, the humped branch $a > 0$, a positive Riccati loading $B'(\tau) > 0$ and a physical invariant distribution with finite mean $\theta_P > 0$. The equilibrium forward curve is nondecreasing at every maturity if and only if*

$$\boxed{\theta_P \leq \frac{m}{\gamma}, \qquad \gamma = \sqrt{a^2 + 2\beta}.} \tag{9.19}$$

Proof. From equation (9.16), monotonicity for every maturity is equivalent to

$$m + a\theta_P - \beta\theta_P B(\tau) \geq 0 \quad \text{for all } \tau. \tag{9.20}$$

Because $B(\tau)$ increases to B_∞, the smallest possible value of the bracket is its long-maturity limit:

$$m + a\theta_P - \beta\theta_P B_\infty = m + [a - (a+\gamma)]\theta_P \tag{9.21}$$

$$= m - \gamma\theta_P. \tag{9.22}$$

It is nonnegative precisely when $\theta_P \leq m/\gamma$. If equality holds, the slope is strictly positive for every finite maturity and approaches zero at infinity. □

The criterion constrains only the *mean* of the physical invariant distribution. It does not force the invariant density to be Gamma and does not uniquely identify b_P or the market price of risk.

9.5 Risk compensation and the equilibrium forward spread

Chapter 7 provided the pointwise identity for the instantaneous expected excess holding-return *rate* of a zero-coupon bond:

$$\mathrm{ERP}_\tau(r) = B(\tau)\big[b_Q(r) - b_P(r)\big], \qquad b_Q(r) = m + ar. \tag{9.23}$$

We now connect that relationship directly to the entire mean forward curve. Taking the stationary physical expectation and using $\mathbb{E}_{\pi_P}[b_P] = 0$,

$$\mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau(r)] = B(\tau)\big[m + a\mathbb{E}_{\pi_P}[r] - \mathbb{E}_{\pi_P}[b_P(r)]\big] \tag{9.24}$$

$$= B(\tau)(m + a\theta_P). \tag{9.25}$$

Thus

$$\boxed{\mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau(r)] = B(\tau)(m + a\theta_P).} \tag{9.26}$$

The quantity on the left is the stationary mean *instantaneous conditional expected excess holding-return rate*, not a finite-horizon expected bond return or a conventional yield term premium.

Now use the Riccati equation to expand the equilibrium curve:

$$x_e^P(\tau) - \theta_P = mB(\tau) + \theta_P[B'(\tau) - 1] \tag{9.27}$$

$$= mB(\tau) + \theta_P\left[aB(\tau) - \frac{\beta}{2}B(\tau)^2\right] \tag{9.28}$$

$$= B(\tau)(m + a\theta_P) - \frac{\beta\theta_P}{2}B(\tau)^2. \tag{9.29}$$

Using equation (9.26), we obtain

$$\boxed{x_e^P(\tau) - \theta_P = \mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau(r)] - \frac{\beta\theta_P}{2}B(\tau)^2.} \tag{9.30}$$

The equilibrium forward spread is the mean instantaneous expected bond risk compensation minus an exact variance-related pricing correction. The second term is not a separately estimated premium and should not be equated automatically with a textbook convexity adjustment for a finite-horizon yield term premium.

Differentiating gives

$$\boxed{\partial_\tau x_e^P(\tau) = \partial_\tau \mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau(r)] - \beta\theta_P B(\tau)B'(\tau).} \tag{9.31}$$

Under our *optional* monotonicity selection, the maturity derivative of expected instantaneous risk compensation must cover the correction $\beta\theta_P BB'$. This is a reformulation of the same criterion, not a second independent restriction.

At the short boundary, equation (9.26) gives

$$\partial_\tau \mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau]|_{\tau=0} = m + a\theta_P = \partial_\tau x_e^P(0), \tag{9.32}$$

because $B(0) = 0$ and $B'(0) = 1$. The mean ERP *level* itself is zero at $\tau = 0$; its *maturity derivative* equals the initial equilibrium-curve slope. Equivalently, since $\sigma\lambda = b_Q - b_P$,

$$\boxed{\partial_\tau x_e^P(0) = \mathbb{E}_{\pi_P}[\sigma(r)\lambda(r)] = m + a\theta_P.} \tag{9.33}$$

This connects the equilibrium geometry to the measure-change mechanism developed in Chapters 6 and 7.

9.6 A preview of the long-forward measure

The criterion $\theta_P \leq m/\gamma$ has a deeper interpretation, which Chapter 10 will derive fully. We introduce it here only to show how the equilibrium calculation points beyond the physical measure.

The long-maturity pricing loading B_∞ determines the decreasing positive pricing eigenfunction $\phi(r) = e^{-B_\infty r}$. Its associated long-forward (recurrent eigenfunction) measure, written $\mathbb{L}$, has short-rate drift

$$b_L(r) = m + (a - \beta B_\infty)r = m - \gamma r. \tag{9.34}$$

The invariant mean under this *different* measure is

$$\boxed{\theta_L = \frac{m}{\gamma}.} \tag{9.35}$$

We will justify the eigenfunction change of measure, its martingale property and the recurrent invariant law in Chapter 10. At this stage, the identity $m/\gamma = \theta_L$ should be read as a preview, not as an independent assumption.

Our monotonicity criterion therefore becomes

$$\boxed{\partial_\tau x_e^P(\tau) \geq 0 \ \forall\tau \quad \Longleftrightarrow \quad \theta_P \leq \theta_L.} \tag{9.36}$$

The physical distribution π_P and the long-forward distribution need not coincide. Even equality of their means would not imply equality of their laws.

One may also rewrite equation (9.16) as

$$\partial_\tau x_e^P(\tau) = B'(\tau)\big[m + (a - \beta B(\tau))\theta_P\big]. \tag{9.37}$$

The bracket is the *initial-date short-rate drift* of the corresponding finite-maturity forward measure, evaluated at $r = \theta_P$. Its limit is $m - \gamma\theta_P$. The coincidence is an exact consequence of the same Riccati pricing geometry; the finite-maturity forward measure is time-inhomogeneous and must not be confused with a new stationary physical law.

9.7 A quantitative consequence of the volatility hump

The optional monotonicity condition becomes particularly restrictive when combined with the humped pricing branch $a > 0$. Recall from Chapter 5 the instantaneous-forward volatility hump ratio

$$R = 1 + \frac{a^2}{2\beta} > 1. \tag{9.38}$$

Introduce the physically meaningful dimensionless ratio

$$u := \frac{\theta_P}{\rho}, \qquad \rho = mB_\infty. \tag{9.39}$$

To express the mean bound in these quantities, first note that $\gamma^2 - a^2 = 2\beta$, hence

$$B_\infty = \frac{a+\gamma}{\beta} = \frac{2}{\gamma - a}, \qquad \rho = \frac{2m}{\gamma - a}. \tag{9.40}$$

Therefore,

$$\frac{\theta_L}{\rho} = \frac{m/\gamma}{2m/(\gamma - a)} = \frac{1}{2}\left(1 - \frac{a}{\gamma}\right). \tag{9.41}$$

Finally, $R = \gamma^2/(2\beta)$ and $a > 0$ imply

$$\frac{a}{\gamma} = \sqrt{1 - \frac{1}{R}}. \tag{9.42}$$

Combining these results with $\theta_P \le \theta_L$ yields

$$\boxed{u = \frac{\theta_P}{\rho} \le \frac{1}{2}\left(1 - \sqrt{1 - \frac{1}{R}}\right) < \frac{1}{2}.} \tag{9.43}$$

For $0 < u < 1/2$, rearranging the inequality equivalently gives

$$\boxed{1 < R \le \frac{1}{4u(1-u)}.} \tag{9.44}$$

The strict $u < 1/2$ restriction comes from combining *global monotonicity* with the *humped branch*. A volatility hump by itself places no such bound on θ_P. The use of R assumes the model's *instantaneous-forward* volatility, not yield volatility or an empirically smoothed proxy.

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

Equation (9.43) is a severe conditional model restriction, not an established regularity of observed Treasury curves. A failure in suitable data would challenge the *joint* assumptions (pricing class, correctly measured hump and global monotonicity); it would not, on its own, disprove physical stationarity or the possibility of humped pricing volatility.

9.8 Two physical CIR examples: same pricing hump, different equilibrium geometry

Consider the pricing parameters $m = a = \beta = 1$. Then $\gamma = \sqrt{3}$, $B_\infty = 1 + \sqrt{3}$, and the volatility hump ratio is $R = 3/2$. Under the physical measure, take the *illustrative* restoring CIR drift

$$b_P(r) = 1 - \kappa_P r, \qquad \kappa_P > 0. \tag{9.45}$$

This has stationary mean $\theta_P = 1/\kappa_P$ and, with the given diffusion variance, a Gamma stationary distribution.

If $\kappa_P = 3$, then $\theta_P = 1/3 < 1/\sqrt{3} = \theta_L$. The equilibrium forward curve rises at every finite maturity and approaches $\rho = 1 + \sqrt{3}$. Its initial slope is $1 + 1/3 = 4/3$.

If instead $\kappa_P = 1$, then $\theta_P = 1 > \theta_L$. The physical diffusion remains stationary, and the *risk-neutral* pricing dynamics and forward-volatility hump remain exactly the same. Yet its equilibrium forward curve rises initially, reaches a maximum and then approaches the same ρ from above.

> The difference between these two cases is not pricing arbitrage or the presence of a volatility hump. It is whether we additionally select a globally nondecreasing *physical stationary mean* forward curve. That is why this geometric requirement must never be silently folded into the definition of equilibrium.

9.9 What this chapter establishes, and what remains free

The stationary averaging calculation and the Riccati equation establish the entire equilibrium forward curve once the scalar physical mean θ_P and pricing parameters are known. They also give its initial slope, long limit, and decomposition into stationary expected bond excess returns and the volatility-related correction. No Gamma distribution has been assumed in deriving any of those formulas.

What remains free is the full physical drift b_P, the stationary density beyond its relevant moments, and the state-dependent market price of risk, unless additional information is supplied. A nondecreasing equilibrium curve is an *optional* selection that restricts θ_P; it does not identify those remaining functions. The result is therefore a compatibility classification rather than a complete econometric identification.

Chapter 10 will develop the long-forward measure rigorously. It will explain why the same quantity m/γ arises as the mean of a recurrent long-forward diffusion, and how that representation relates to—but remains distinct from—physical equilibrium.

Chapter 10

Long-Horizon Pricing and the Long-Forward Measure

Chapter 9 showed that global monotonicity of the physical stationary expected forward curve is equivalent to a bound on its mean short rate. In this chapter we explain why the same bound can be expressed using the invariant mean of a long-forward measure*. Beginning with the risk-neutral pricing operator, we construct its positive recurrent eigenfunction, verify the associated stochastic change of measure, and derive the restoring square-root diffusion under the resulting measure. We also distinguish finite-maturity forward measures from the long-forward limit, explain the relevance of eigenfunction selection, and show precisely where the optional physical equilibrium geometry enters. No equality between the physical and long-forward probability laws is assumed.*

10.1 Three different questions, and three different measures

By now we have distinguished physical forecasting from risk-neutral valuation. The physical measure P describes the assumed statistical evolution of the short rate. The risk-neutral measure Q is the measure under which discounted traded bond prices are martingales, subject to the usual market and measure-change hypotheses. We now introduce a third measure, denoted L, which is naturally associated with *very long-horizon* pricing.

These are different questions: what is likely to happen; how is a finite-maturity claim priced; and what stochastic dynamics appear after factoring out the dominant long-horizon component of the pricing operator? The third does not require either $P = Q$ or $P = L$.

We retain the humped pricing branch derived in Chapters 4 and 5:

$$\mathrm{d}r_t = (m + ar_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^Q, \tag{10.1}$$

$$P(\tau, r) = \exp[-A(\tau) - B(\tau)r], \tag{10.2}$$

$$A' = mB, \qquad B' = 1 + aB - \tfrac{1}{2}\beta B^2, \qquad A(0) = B(0) = 0. \tag{10.3}$$

Here $m, \beta > 0$ and $a > 0$; we assume parameters and boundary behaviour making the square-root process nonnegative and the changes of measure below well defined. The standard sufficient strict-positivity condition from a positive initial state is $2m \geq \beta$. It is not needed merely to manipulate the Riccati equation.

10.2 The long-maturity Riccati fixed point

The bond loading B increases from zero towards the positive equilibrium of its Riccati equation. Write

$$\gamma := \sqrt{a^2 + 2\beta}, \qquad B_\infty := \frac{a+\gamma}{\beta}, \qquad B_- := \frac{a-\gamma}{\beta} < 0. \tag{10.4}$$

Indeed, the fixed-point equation is

$$0 = 1 + aB - \tfrac{1}{2}\beta B^2 = -\tfrac{1}{2}\beta(B - B_\infty)(B - B_-). \tag{10.5}$$

Since $B(0) = 0$ lies between the two roots, $B'(\tau) > 0$ until the positive fixed point is reached asymptotically. Thus

$$B(\tau) \uparrow B_\infty, \qquad B'(\tau) \longrightarrow 0. \tag{10.6}$$

In particular, $A'(\tau) = mB(\tau)$ approaches mB_∞. Defining

$$\boxed{\rho := mB_\infty = \frac{m(a+\gamma)}{\beta}}, \tag{10.7}$$

we obtain the limiting instantaneous forward rate

$$x(\tau, r) = mB(\tau) + B'(\tau)r \longrightarrow \rho. \tag{10.8}$$

This limit is independent of the current short-rate state. It is a long-maturity *forward rate*, not a prediction that the physical short rate converges to ρ.

The risk-neutral conditional short-rate mean need not converge. Indeed $\mathbb{E}_r^Q[r_t] = (r + m/a)e^{at} - m/a$ grows when $a > 0$, while the Riccati bond loading approaches a finite fixed point. A pricing expectation of an exponential discount factor is not the exponential of an expected short-rate path.

10.3 Why a pricing eigenfunction appears

The risk-neutral generator acts on a smooth test function g as

$$\mathcal{L}_Q g(r) = (m + ar)g'(r) + \tfrac{1}{2}\beta r g''(r). \tag{10.9}$$

For interest-rate pricing the relevant operator is not $\mathcal{L}_Q$ alone: future cash flows are discounted at the short rate. Define the *killed pricing operator*

$$\mathcal{P}g(r) := (\mathcal{L}_Q - r)g(r). \tag{10.10}$$

We seek a positive function ϕ satisfying

$$\boxed{(\mathcal{L}_Q - r)\phi = -\rho\phi.} \tag{10.11}$$

This is analogous to a familiar eigenvalue equation in differential equations. The function ϕ captures how the dominant long-horizon pricing component depends on the current state; ρ is the associated long-run decay rate of discounted values.

Guided by the exponential-affine bond-price representation, set

$$\phi(r) = e^{-B_\infty r}. \tag{10.12}$$

Then $\phi' = -B_\infty \phi$ and $\phi'' = B_\infty^2 \phi$. Substituting into the pricing operator gives

$$(\mathcal{L}_Q - r)\phi = \left[-mB_\infty + r\left(-aB_\infty + frac12\beta B_\infty^2 - 1\right)\right]\phi \tag{10.13}$$

$$= -mB_\infty \phi = -\rho\phi, \tag{10.14}$$

where the coefficient of r vanishes by the Riccati fixed-point equation. This verifies (10.11) directly.

A positive solution of the eigenvalue equation does not automatically determine the economically relevant long-forward measure. Different positive eigenfunctions may exist, and a candidate stochastic density must be a *true martingale*. We will identify the solution associated with recurrent transformed dynamics in this explicit model.

10.4 From an eigenfunction to a change of probability measure

Consider the process

$$\boxed{M_t = \exp\left(\rho t - \int_0^t r_s \, \mathrm{d}s\right) \frac{\phi(r_t)}{\phi(r_0)}.} \tag{10.15}$$

The factor $\exp(-\int_0^t r_s \, \mathrm{d}s)$ is the stochastic discount factor under Q. The factor $e^{\rho t}$ removes the dominant long-horizon exponential decay, and the ratio $\phi(r_t)/\phi(r_0)$ adjusts for the current and future states.

To see how M evolves, apply Itô's formula to $\phi(r_t)$. Since

$$\mathrm{d}\phi(r_t) = \mathcal{L}_Q \phi(r_t) \, \mathrm{d}t + \phi'(r_t)\sqrt{\beta r_t} \, \mathrm{d}W_t^Q, \tag{10.16}$$

and $\mathcal{L}_Q \phi = (r - \rho)\phi$, the drift cancels against the differential of the exponential factor in (10.15). Thus

$$\boxed{\frac{\mathrm{d}M_t}{M_t} = -B_\infty \sqrt{\beta r_t} \, \mathrm{d}W_t^Q.} \tag{10.17}$$

So M is a positive local martingale. For a valid probability change we require $\mathbb{E}_r^Q[M_t] = 1$, not merely the vanishing formal drift.

In this model the expectation can be checked without relying on a general exponential-martingale criterion. For fixed t, the affine pricing transform gives

$$\mathbb{E}_r^Q\left[e^{-\int_0^t r_s \, \mathrm{d}s} e^{-B_\infty r_t}\right] = e^{-\rho t - B_\infty r}. \tag{10.18}$$

The reason is that a transform starting with terminal loading B_∞ remains at the Riccati fixed point, and the associated constant term grows at rate mB_∞. The bounded payoff $0 < e^{-B_\infty r_t} \leq 1$, together with standard Feynman–Kac uniqueness for this nonnegative square-root model, validates the expectation identity. Substituting (10.18) into (10.15) yields

$$\boxed{\mathbb{E}_r^Q[M_t] = 1.} \tag{10.19}$$

The process is therefore a true martingale on each finite horizon, and we may define L by

$$\left.\frac{\mathrm{d}L}{\mathrm{d}Q}\right|_{\mathcal{F}_t} = M_t. \tag{10.20}$$

The finite-horizon laws are equivalent under the stated positivity and regularity assumptions. Equivalence on every finite horizon should not be confused with equivalence on an infinite time interval.

10.5 The restoring short-rate dynamics under the long-forward measure

Equation (10.17) gives the Brownian change of measure. Under L,

$$\mathrm{d}W_t^L = \mathrm{d}W_t^Q + B_\infty \sqrt{\beta r_t}\,\mathrm{d}t \tag{10.21}$$

is Brownian motion. Substituting $\mathrm{d}W_t^Q = \mathrm{d}W_t^L - B_\infty\sqrt{\beta r_t}\,\mathrm{d}t$ into (10.1) yields

$$\mathrm{d}r_t = [m + ar_t - \beta B_\infty r_t]\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^L \tag{10.22}$$

$$= \boxed{(m - \gamma r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^L}, \tag{10.23}$$

because $\beta B_\infty = a + \gamma$.

We have arrived at a notable contrast: the same pricing model has a *non-restoring* drift $m + ar$ under Q, but a *restoring* drift $m - \gamma r$ under L. The long-forward measure is not the physical measure. Its mean-reversion does not prove physical mean reversion; the latter must be established separately.

The L dynamics are CIR with speed γ and long-run mean

$$\boxed{\theta_L = \frac{m}{\gamma}.} \tag{10.24}$$

Solving the stationary zero-current equation from Chapter 8 gives

$$\boxed{\pi_L(r) = \frac{(2\gamma/\beta)^{2m/\beta}}{\Gamma(2m/\beta)} r^{2m/\beta - 1} e^{-2\gamma r/\beta}, \qquad r > 0.} \tag{10.25}$$

The L invariant density is Gamma. This *does not* imply that the physical invariant density π_P is Gamma; Chapter 8 supplied explicit non-Gamma physical examples compatible with the same pricing specification.

10.6 Why the recurrent eigenfunction matters

The Riccati fixed-point equation has another root, $B_- < 0$. The positive function $e^{-B_- r}$ formally satisfies a pricing eigenfunction equation too, with eigenvalue mB_-. Its associated drift is

$$m + (a - \beta B_-)r = m + \gamma r, \tag{10.26}$$

which is non-restoring. It does not give a recurrent transformed short-rate process. Moreover, it grows exponentially in r, so one must not assume its putative density process is a true martingale without a separate check.

For the long-forward construction used here, the decreasing eigenfunction $e^{-B_\infty r}$ produces a true martingale and recurrent transformed diffusion. This is the relevant selection in the present affine model. The broader theory of recurrent eigenfunctions and long-term factorisation is developed, for example, by Hansen and Scheinkman and by Qin and Linetsky.

10.7 Finite-maturity forward measures and the limiting drift

A long-forward measure is easier to understand if we first consider the ordinary forward measure for a fixed maturity date T. Taking the T-bond as numeraire leads to a measure Q^T under which

its appropriately normalised price is the pricing density. At calendar time $s < T$ the short-rate drift is

$$b_{Q^T}(r; s) = b_Q(r) + \beta r\, \partial_r \log P(T - s, r). \tag{10.27}$$

Since $\partial_r \log P(\tau, r) = -B(\tau)$, this becomes

$$\boxed{b_{Q^T}(r; s) = m + [a - \beta B(T - s)]\, r.} \tag{10.28}$$

This drift is *time dependent*: the remaining maturity $T - s$ changes as calendar time passes. It would be incorrect to describe Q^T as a new time-homogeneous stationary model at finite T.

For any fixed observation time s, the drift in (10.28) converges as $T \to \infty$ to

$$\lim_{T\to\infty} b_{Q^T}(r; s) = m - \gamma r = b_L(r). \tag{10.29}$$

This illustrates the approach to the long-forward measure. More precisely, the finite-maturity bond-numeraire density tends to the eigenfunction density (10.15) on each fixed finite observation horizon, because $B(T) - B(T - t) \to 0$ and $A(T) - A(T - t) \to \rho t$; the bond-price ratio $P(T - t, r_t)/P(T, r_0)$ converges to $e^{\rho t}\phi(r_t)/\phi(r_0)$ in the associated numeraire formula.

10.8 The equilibrium-curve slope and the long-forward mean

Chapter 9 established that, if the physical diffusion has an invariant probability distribution with finite mean θ_P, its stationary mean instantaneous-forward curve is

$$x_e^P(\tau) = mB(\tau) + \theta_P B'(\tau). \tag{10.30}$$

Differentiating and using $B'' = (a - \beta B)B'$ gives

$$\partial_\tau x_e^P(\tau) = mB'(\tau) + \theta_P B''(\tau) \tag{10.31}$$

$$= B'(\tau)\big[m + (a - \beta B(\tau))\theta_P\big]. \tag{10.32}$$

Now compare the bracket with (10.28), at the initial date for a bond maturing at $T = \tau$:

$$\boxed{\partial_\tau x_e^P(\tau) = B'(\tau)\, b_{Q^\tau}(\theta_P; 0).} \tag{10.33}$$

The notation means *evaluate the finite-maturity forward-measure drift function at the numerical value* $r = \theta_P$. It does not mean average that measure's drift using its own invariant distribution.

The right side is the product of $B'(\tau) > 0$ and a bracket that decreases with maturity:

$$b_{Q^\tau}(\theta_P; 0) = m + (a - \beta B(\tau))\theta_P \downarrow m - \gamma\theta_P. \tag{10.34}$$

Therefore the slope is nonnegative at every finite maturity exactly when this limiting bracket is nonnegative. Thus

Proposition 10.1 (Long-forward interpretation of global monotonicity)**.** *Assume the humped affine pricing model, an admissible physical stationary law with* $0 < \theta_P < \infty$*, and the differentiability conditions above. Then*

$$\boxed{\partial_\tau x_e^P(\tau) \geq 0 \quad \textit{for all } \tau \geq 0 \quad \Longleftrightarrow \quad \theta_P \leq \theta_L = \frac{m}{\gamma}.} \tag{10.35}$$

Proof. Equation (10.32) has a strictly positive multiplier B'. Its bracket decreases from $m + a\theta_P$ to $m - \gamma\theta_P$. Hence it stays nonnegative at all maturities if and only if its limit is nonnegative. Dividing by $\gamma > 0$ gives (10.35). □

The comparison $\theta_P \le \theta_L$ is *not* a general requirement of no-arbitrage, of physical stationarity, or of the volatility hump. It is equivalent to the *additional* choice that the stationary physical expected forward curve be globally nondecreasing. A physically stationary model with $\theta_P > \theta_L$ is compatible with the same humped pricing model, but its mean forward curve eventually declines towards its long limit.

10.9 The point of maximum forward volatility is not the point of maximum equilibrium forwards

There is a subtle distinction here that is worth making explicit. The volatility loading $B'(\tau)$ reaches its hump at the maturity satisfying

$$B(\tau_*) = \frac{a}{\beta}. \tag{10.36}$$

At that maturity $B''(\tau_*) = 0$, whereas equation (10.32) gives

$$\partial_\tau x_e^P(\tau_*) = mB'(\tau_*) > 0. \tag{10.37}$$

Thus the stationary expected forward curve is *still increasing* at the maturity of the instantaneous-forward volatility maximum. If $\theta_P > \theta_L$, the expected forward curve reaches its own maximum later, when

$$B(\tau_{e,*}) = \frac{a + m/\theta_P}{\beta} > \frac{a}{\beta}. \tag{10.38}$$

This turning point is attained at finite maturity precisely when $\theta_P > m/\gamma$, because $B(\tau)$ rises only as far as $B_\infty = (a + \gamma)/\beta$. When $\theta_P \le \theta_L$, no finite-maturity equilibrium turning point exists.

10.10 Changes of measure and what remains unidentified

Under our sign convention the physical and pricing drifts satisfy

$$b_Q(r) = b_P(r) + \sqrt{\beta r}\,\lambda(r). \tag{10.39}$$

The long-forward drift follows from the eigenfunction transform:

$$b_L(r) = b_Q(r) - \beta B_\infty r = m - \gamma r. \tag{10.40}$$

Therefore the physical-to-long-forward risk adjustment is

$$\boxed{\lambda_L(r) := \frac{b_L(r) - b_P(r)}{\sqrt{\beta r}} = \lambda(r) - B_\infty\sqrt{\beta r}.} \tag{10.41}$$

This is an identity between drift differences, not a declaration that any of the three measures are equal. In particular, $\lambda_L = 0$ would imply $b_P = b_L$; we impose no such condition. The existence and integrability of an actual P-to-L density process still require the appropriate change-of-measure hypotheses.

The long-forward measure is fixed by the selected Q pricing parameters in this model. The physical stationary law remains free within the admissible class, subject to any optional mean restriction. Thus the long-forward mean gives a *reference value* for a geometric selection without identifying the physical drift, its stationary density, or its market price of risk.

10.11 What this chapter establishes, and what remains free

We have constructed the long-forward measure rather than simply naming it. Starting from the Riccati fixed point, we verified a pricing eigenfunction, exhibited a martingale density, derived its Girsanov drift adjustment, and obtained a restoring CIR diffusion with mean $\theta_L = m/\gamma$. We also showed how finite-maturity forward-measure drifts approach the long-forward drift and why the same coefficient determines the slope of the physical stationary mean forward curve.

The resulting mean ordering $\theta_P \leq \theta_L$ remains conditional on the additional global monotonicity assumption. Neither the existence of a long-forward measure nor the volatility hump selects a unique physical stationary distribution. In the next part of the monograph we can therefore ask which ingredients of this compatibility hierarchy are identifiable in data, and which require independent physical forecasts or bond-return information.

Chapter 11

What Compatibility Determines, and What It Does Not

The preceding chapters have established a hierarchy of conditional relationships linking physical forecasts, affine pricing, forward-volatility geometry, bond risk compensation, stationary physical distributions and long-forward asymptotics. This chapter asks a different question: which objects does that hierarchy actually determine? We separate a model's internal compatibility from identification using financial observations. Two constructions show why even exact knowledge of the risk-neutral pricing curve does not select a unique physical equilibrium. Conversely, the maturity structure of conditional bond excess returns can, in principle, recover the physical drift. We catalogue observable restrictions, explain the limits of stationary-mean and long-forward information, and specify what additional data would be required to distinguish competing physical models. The objective is neither to extend the model by new assumptions nor to claim uniqueness where only coexistence has been proved.

11.1 From coexistence to identification

The first ten chapters addressed whether several mathematical descriptions can hold simultaneously. We began with a one-state short rate, selected an affine state-to-curve map through exact forecast closure, imposed risk-neutral bond pricing and a zero-variance boundary, and studied the resulting square-root diffusion. The volatility hump selected a positive coefficient in the risk-neutral drift. The physical diffusion could nonetheless be restoring and stationary, because its drift differs from the risk-neutral drift by the market price of risk.

The existence of such a coherent system is a *compatibility* result. It does not automatically tell us whether the system's unknown parameters or functions are uniquely recoverable from observations. That further question is one of *identification*. Two distinct compatible models can produce the same observed bond prices; if so, pricing observations alone cannot discriminate between them.

For clarity, we retain the humped one-state pricing model:

$$\mathrm{d}r_t = b_P(r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\;\mathrm{d}W_t^P, \tag{11.1}$$

$$\mathrm{d}r_t = (m + a r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\;\mathrm{d}W_t^Q, \tag{11.2}$$

$$P(\tau, r) = \exp[-A(\tau) - B(\tau) r], \tag{11.3}$$

$$A' = mB, \qquad B' = 1 + aB - \tfrac{1}{2}\beta B^2. \tag{11.4}$$

Here $m, a, \beta > 0$ and $A(0) = B(0) = 0$. Throughout, the chosen physical process is assumed to be an admissible conservative diffusion on $(0, \infty)$, with an invariant law where stationarity is invoked, and the relevant P–Q measure changes are assumed valid on each finite horizon. These are substantive hypotheses, not consequences of algebraic drift matching.

The central distinction. Compatibility asks whether a physical model and a pricing model can describe the same state process under different measures. Identification asks which observations determine which model ingredients, and with what uniqueness. A proof of the first is not a proof of the second.

11.2 The hierarchy, revisited without compressing its assumptions

It is useful to distinguish the mathematical selections from their consequences. The following is an ordered construction, not a sequence in which each line is a theorem following automatically from the previous line.

1. **A one-state representation.** The forward curve is a function $x(\tau, r)$ of remaining maturity and the current short rate. This alone does not make x affine.
2. **Exact physical forecast closure.** The additional condition

$$\mathbb{E}_r^P[x(\tau, r_h)] = x(\tau, \mathbb{E}_r^P[r_h]) \tag{11.5}$$

   for all eligible states, maturities and short horizons selects affinity in r, given adequate smoothness and nonzero diffusion variance. It is not required by one-state Markov modelling or no-arbitrage.
3. **Arbitrage-free pricing.** Exponential-affine bond prices must satisfy the Q pricing PDE. With sufficiently distinct nonzero maturity loadings, the master equation selects affine Q drift and affine instantaneous variance.
4. **Positive square-root boundary.** Selecting zero instantaneous variance at $r = 0$ removes the constant variance term and gives $\sigma^2(r) = \beta r$. Boundary admissibility is examined separately.
5. **An interior instantaneous-forward volatility hump.** Within this branch the hump occurs exactly when $a > 0$; it is not a prediction that all term structures must be humped.
6. **Physical stationarity.** The chosen b_P must support an invariant law π_P with the appropriate boundary, recurrence and integrability conditions. This does not require b_P to be CIR.
7. **Optional global equilibrium-curve monotonicity.** Requiring $\partial_\tau x_e^P(\tau) \geq 0$ for all maturities places a bound on the stationary physical mean; it is not implied by stationarity.

This is a *hierarchy of conditional compatibility*, not an attempt to establish every item as a universal financial law. The strength of the resulting conclusions must be judged together with the strength of the selected hypotheses.

11.3 What pricing geometry determines

Because the bond price is exponential-affine, the instantaneous forward rate and its diffusion loading are

$$x(\tau, r) = mB(\tau) + B'(\tau)r, \tag{11.6}$$

$$q(\tau, r) = \sqrt{\beta r}\, B'(\tau). \tag{11.7}$$

The Riccati equation therefore fixes the relative maturity dependence of the forward-rate volatility, independently of b_P.

Differentiating gives

$$B''(\tau) = [a - \beta B(\tau)]B'(\tau). \tag{11.8}$$

In the humped branch $a > 0$, $B(0) = 0$ and $B'(\tau) > 0$. Thus the unique interior maximum occurs where

$$B(\tau_*) = \frac{a}{\beta}. \tag{11.9}$$

Write $\gamma = \sqrt{a^2 + 2\beta}$. The hump's relative height and location satisfy

$$R := \frac{B'(\tau_*)}{B'(0)} = 1 + \frac{a^2}{2\beta}, \tag{11.10}$$

$$\tau_* = \frac{1}{\gamma} \log \frac{\gamma + a}{\gamma - a}. \tag{11.11}$$

These give a useful structural identification exercise. Defining $z = a/\gamma$, we have

$$z = \sqrt{1 - \frac{1}{R}}, \qquad \gamma = \frac{2\,\mathrm{artanh}(z)}{\tau_*}, \qquad a = \gamma z, \qquad \beta = \tfrac{1}{2}\gamma^2(1 - z^2). \tag{11.12}$$

Thus, if the model is exactly correct and the *instantaneous-forward volatility loading* is observed without error, the hump ratio and its maturity identify a and β. The intercept m can then be identified from the forward-curve level, for instance its long-maturity limit $\rho = m(a + \gamma)/\beta$.

Structural versus empirical identification. Equation (11.12) is an inversion of an idealised model. Forward curves are inferred from noisy bond prices; forward volatilities require time-series estimation and often differentiation of fitted yields. Observing a hump in some fitted volatility series is not, on its own, evidence that the true loading is the Riccati loading B', nor that the branch $a > 0$ is uniquely identified.

11.4 What pricing geometry cannot determine: two physical laws

Once the Q parameters are fixed, the risk-premium identity reads

$$\sqrt{\beta r}\,\lambda(r) = m + ar - b_P(r). \tag{11.13}$$

We may choose different admissible physical drifts and then use this identity to construct corresponding market prices of risk. This is how the pricing geometry can coexist with more than one physical stationary distribution.

11.4.1 A restoring CIR physical law

Choose $\kappa > 0$ and set

$$b_P^{(1)}(r) = m - \kappa r. \tag{11.14}$$

When the relevant boundary conditions hold, its stationary density is Gamma:

$$\pi_1(r) = C_1 r^{2m/\beta - 1} \exp\left(-\frac{2\kappa}{\beta} r\right), \qquad \theta_1 = \frac{m}{\kappa}. \tag{11.15}$$

The associated risk-price function is

$$\lambda_1(r) = \frac{a + \kappa}{\sqrt{\beta}} \sqrt{r}. \tag{11.16}$$

11.4.2 A nonlinear restoring physical law

Alternatively, choose $\delta > 0$ and define

$$b_P^{(2)}(r) = m - \kappa r - \delta \frac{r}{1+r}. \tag{11.17}$$

The zero-current stationary density takes the form

$$\pi_2(r) = C_2 r^{2m/\beta - 1} \exp\left(-\frac{2\kappa r}{\beta}\right)(1+r)^{-2\delta/\beta}. \tag{11.18}$$

For $\kappa, \delta > 0$ this density is integrable at infinity and, with $m > 0$, locally integrable near zero; further boundary conditions select a conservative positive recurrent diffusion. The risk price is now

$$\lambda_2(r) = \frac{(a+\kappa)r + \delta r/(1+r)}{\sqrt{\beta r}}. \tag{11.19}$$

Both laws share the *same* Q drift, the *same* bond prices and the *same* volatility hump, although their physical drifts and invariant laws differ. In either case the change of measure must be justified, not inferred solely from its formal drift expression.

The two physical means need not coincide. Indeed, holding m, κ, β fixed, the extra negative term in (11.17) shifts the drift and generally changes the stationary mean. But even knowledge of the stationary mean does not resolve the nonuniqueness: distinct admissible densities can be constructed with the same first moment, while their logarithmic derivatives and physical drifts differ.

Compatibility is not uniqueness. The volatility hump fixes an aspect of the Q model. It does not pick a unique P model, even when physical stationarity is imposed. The market price of risk changes together with the physical law so that the risk-neutral pricing equations remain valid.

11.5 What additional information can identify the physical drift?

The central result of Chapter 7 supplies an answer. Let $\mathrm{ERP}_\tau(r)$ denote the instantaneous *conditional expected excess holding-return rate* for a zero-coupon bond with remaining maturity τ. Then

$$\mathrm{ERP}_\tau(r) = B(\tau)[b_Q(r) - b_P(r)] = B(\tau)[m + ar - b_P(r)]. \tag{11.20}$$

At any maturity with $B(\tau) > 0$, this can be inverted:

$$\boxed{b_P(r) = m + ar - \frac{\mathrm{ERP}_\tau(r)}{B(\tau)}.} \tag{11.21}$$

Thus exact knowledge of the pricing parameters *and* the state-dependent conditional instantaneous ERP identifies the physical drift within the model. It does not require a prior assumption of CIR or Gamma stationarity.

The same reasoning makes the maturity restriction explicit. For all positive maturities,

$$\frac{\mathrm{ERP}_{\tau_1}(r)}{B(\tau_1)} = \frac{\mathrm{ERP}_{\tau_2}(r)}{B(\tau_2)} = m + ar - b_P(r). \tag{11.22}$$

Failure of exact factorisation would contradict the selected one-state model, assuming that the corresponding expected instantaneous returns and bond sensitivities have been correctly estimated. The relation is standard affine pricing in a revealing compatibility form, not a new independent no-arbitrage law.

11.5.1 From physical drift to stationary density

Under zero stationary probability current, with $\sigma^2(r) = \beta r$, the Fokker–Planck equation gives

$$\pi_P(r) = \frac{C}{\beta r} \exp\left(\int^r \frac{2b_P(u)}{\beta u}\, \mathrm{d}u \right). \tag{11.23}$$

Substituting (11.21) shows the inverse link:

$$\boxed{\pi_P(r) = C r^{2m/\beta - 1} \exp\left[\frac{2ar}{\beta} - \frac{2}{\beta} \int^r \frac{\mathrm{ERP}_\tau(u)}{B(\tau)u}\, \mathrm{d}u \right].} \tag{11.24}$$

The normalising constant absorbs reference-point factors. This yields a *candidate* invariant density; one must still verify normalisation, boundary inaccessibility or the selected boundary extension, recurrence and required moments. The right-hand side must be independent of the maturity chosen, if the one-state model holds exactly.

11.5.2 Forecast evidence offers a second route

The physical drift is also the initial slope of the conditional physical forecast:

$$b_P(r) = \frac{\partial}{\partial h} \mathbb{E}_r^P[r_h]\Big|_{h=0}. \tag{11.25}$$

If that derivative were observed or consistently estimated at many states, it could be compared directly with the drift inferred from bond excess returns. Their equality is a further cross-description consistency check. In reality neither a continuous-state conditional drift nor an instantaneous ERP curve is observed without statistical error.

11.6 Why equilibrium-curve information alone is not enough

Physical stationarity gives $\theta_P = \mathbb{E}_{\pi_P}[r]$, and the affine equilibrium forward curve is

$$x_e^P(\tau) = mB(\tau) + \theta_P B'(\tau). \tag{11.26}$$

Every stationary law with the same mean θ_P produces the same curve (11.26), irrespective of higher moments or the local shape of b_P. This is a sharp statement of non-identification, not merely a practical difficulty.

The equilibrium risk-premium expectation is similarly reduced to the mean:

$$\mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau] = B(\tau)(m + a\theta_P), \tag{11.27}$$

where we have used $\mathbb{E}_{\pi_P}[b_P] = 0$ and assumed the required integrability. Even an exactly known *stationary mean* ERP curve cannot distinguish physical laws having the same θ_P.

For the optional condition that x_e^P be nondecreasing at every maturity, Chapters 9 and 10 gave

$$\theta_P \le \theta_L = \frac{m}{\gamma}, \qquad \gamma = \sqrt{a^2 + 2\beta}. \tag{11.28}$$

This reduces the admissible physical means to an interval, but it does not specify a density within that interval. The long-forward law itself is a particular restoring CIR law under L, not the physical law under P.

> A restriction on one moment should not be mistaken for a restriction on an entire probability density. In particular, $\theta_P \le \theta_L$ does not imply $\pi_P = \pi_L$, and it does not require $b_P(r) = m - \gamma r$.

11.7 What the model makes observable and falsifiable

Our assumptions are strong enough to generate restrictions that can, at least conceptually, be contradicted by data. They should be stated separately from empirical demonstrations of the model.

11.7.1 Cross-tenor rank one

At a fixed state, the instantaneous-forward diffusion loadings are proportional to the same Brownian shock:

$$\mathrm{d}x_t(\tau)\big|_{\text{martingale}} = B'(\tau)\sqrt{\beta r_t}\ \mathrm{d}W_t^P. \tag{11.29}$$

Hence the instantaneous conditional covariance kernel is

$$\frac{\mathrm{Cov}_t(\,\mathrm{d}x_t(\tau_i),\,\mathrm{d}x_t(\tau_j))}{\mathrm{d}t} = \beta r_t B'(\tau_i)B'(\tau_j), \tag{11.30}$$

which has rank one across maturities in the exact model. The same property holds for zero-coupon yield innovations, with maturity loading $B(\tau)/\tau$ instead of $B'(\tau)$. A high first principal component share is not the same as *exact* rank one.

11.7.2 The yield loading is more directly accessible

Since the zero-coupon yield is

$$y(\tau, r) = \frac{A(\tau)}{\tau} + \frac{B(\tau)}{\tau} r, \tag{11.31}$$

its state loading is $\ell(\tau) = B(\tau)/\tau$. The Riccati equation imposes the nonlinear relationship

$$\ell(\tau) + \tau\ell'(\tau) = 1 + a\tau\ell(\tau) - \frac{\beta}{2}\tau^2\ell(\tau)^2, \qquad \ell(0+) = 1. \tag{11.32}$$

Thus the model requires more than a dominant common movement in yields: it requires a particular *shape* of the common loading. Tests of the loading and tests of the physical drift are different exercises.

11.7.3 The geometric mean bound

Let $\rho = mB_\infty$ and $u = \theta_P/\rho$. On the humped branch, with R as in (11.10), optional global monotonicity yields

$$0 < u \le \frac{1 - \sqrt{1 - 1/R}}{2} < \frac{1}{2}. \tag{11.33}$$

This is a conditional restriction on physical and long-forward quantities, not a universal bound on observed short and long yields. In particular, using short-maturity yield proxies and fitted long forwards does not automatically produce a test with known rejection properties. One needs reliable estimates, uncertainty bounds and a defensible link between proxies and the true stationary mean.

11.8 A map of what is fixed, free and additionally observable

Object	Determined by	What remains / required evidence
State-to-curve affinity	Exact forecast closure, regular diffusion	Closure itself is a modelling selection
Risk-neutral a, β, m	Ideal full pricing and volatility geometry	Estimation error; branch discrimination in data
Volatility hump R, τ_*	Riccati loading on $a > 0$ branch	Whether true instantaneous loading has this form
Physical drift $b_P(r)$	Q drift plus state-dependent instantaneous ERP	Requires physical forecast or expected-return information
Stationary density π_P	b_P, variance, boundaries, zero current	Recurrence, integrability and boundary classification
Stationary mean θ_P	Physical invariant distribution or its estimate	Does not specify the rest of π_P
Long-forward mean θ_L	Pricing eigenfunction and Q parameters	Not generally equal to θ_P
Monotonicity bound	Optional global shape condition	Rejects some means, not every non-CIR law

One lesson of the table is that an economic observable must be matched with the *mathematical object* it is claimed to identify. An estimated conditional expected holding return is not a yield term premium. A fitted instantaneous-forward volatility hump is not automatically the volatility loading of a one-state diffusion. And an equilibrium average is not a complete physical transition law.

11.9 What this chapter establishes, and what remains free

The compatibility hierarchy unifies several familiar ingredients without claiming that their joint validity uniquely selects the physical dynamics. Exact forecast closure and affine risk-neutral pricing constrain the geometry of the forward curve. The square-root boundary selection and the interior hump further narrow the pricing branch. Under physical stationarity, a broad class of restoring physical laws is still possible. The market price of risk reconciles these with Q; expected instantaneous bond excess returns, if independently identified, would recover the otherwise free physical drift.

Two exact non-identification statements have emerged. First, knowing the risk-neutral bond-price and instantaneous-forward volatility geometry does not determine the physical stationary density. Second, even knowing the entire physical *stationary expected forward curve* identifies only its mean short rate, not the density's shape. The optional long-forward mean inequality narrows that mean but does not remove either ambiguity.

This makes the framework both *restrictive* and *honest about its remaining freedom.* Chapter 12 turns from structural identities to empirical measurement: which required quantities are observable, which must be estimated, and whether the available interest-rate data can discriminate the humped one-state branch from less restrictive alternatives.

Chapter 12

Empirical Measurement and Tests of Compatibility

The preceding chapters derive conditional relationships within a one-state affine term-structure model. This chapter asks what market observations could actually test them. We distinguish four increasingly strong empirical requirements: approximate common movement in yields, the particular Riccati maturity loading, the square-root state-dependent covariance scale, and the joint consistency of physical forecasts, pricing and expected returns. We reproduce the limited historical diagnostics reported in the companion research paper, without claiming to have performed a new calibration. The fitted Treasury curves support useful descriptive comparisons but do not identify the positive-drift risk-neutral branch: a flat loading fits almost as well. Short- and long-tenor averages likewise do not identify the physical stationary mean. We conclude with a sequence of more discriminating tests and explain why a negative empirical finding would not invalidate the conditional compatibility mathematics.

12.1 The purpose of empirical work

The theory of the preceding chapters concerns a carefully specified stochastic system. Under the risk-neutral measure, and within the selected positive square-root branch,

$$\mathrm{d}r_t = (m + ar_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^Q, \qquad m, \beta > 0, \tag{12.1}$$

$$P(\tau, r) = \exp[-A(\tau) - B(\tau)r], \tag{12.2}$$

$$A'(\tau) = mB(\tau), \qquad B'(\tau) = 1 + aB(\tau) - \frac{\beta}{2}B(\tau)^2, \tag{12.3}$$

with $A(0) = B(0) = 0$. The interior instantaneous-forward-volatility hump requires $a > 0$. Under the physical measure, the drift $b_P(r)$ remains a separate object, linked to pricing through the market price of risk. Equilibrium additionally requires an admissible invariant physical law.

It is tempting to ask whether observed yield changes "look like" the functions in these equations. That is a useful beginning, but it is not enough. Market observations are discrete, noisy and often inferred from fitted curves. More importantly, two very different stochastic models may produce similar yield-loading shapes. An empirical exercise must therefore specify which mathematical prediction it is testing, which quantities have been measured rather than assumed, and which alternative explanations have been ruled out.

Empirical hierarchy. An apparently good fit is not a compatibility test until we know which component of the hierarchy generated the fit. Approximate one-factor covariance, Riccati geometry, the pricing measure, physical stationarity and the market price of risk are distinct hypotheses.

12.2 Why yields are easier to observe than instantaneous forwards

Let $y_t(\tau)$ denote the continuously compounded zero-coupon yield and $x_t(\tau)$ the instantaneous forward rate. The accounting identities are

$$y_t(\tau) = -\frac{1}{\tau}\log P_t(\tau) = \frac{1}{\tau}\int_0^\tau x_t(u)\,\mathrm{d}u, \tag{12.4}$$

and, when the maturity dependence is differentiable,

$$x_t(\tau) = \frac{\partial}{\partial\tau}[\tau y_t(\tau)] = y_t(\tau) + \tau\partial_\tau y_t(\tau). \tag{12.5}$$

These expressions are elementary, but their statistical implications are not. An observed bond price generates a yield for a specified cash-flow schedule. By contrast, an instantaneous forward at an arbitrary maturity typically requires interpolation or differentiation of a fitted yield curve. Differentiation magnifies small errors in the fitted maturity profile, particularly at longer tenors where data are sparse and the fit is influenced by its functional form.

Our one-state model suggests a less fragile starting point. From (12.2),

$$y_t(\tau) = \frac{A(\tau)}{\tau} + \frac{B(\tau)}{\tau}r_t. \tag{12.6}$$

Define the *yield loading*

$$\ell(\tau) = \frac{B(\tau)}{\tau}. \tag{12.7}$$

Thus $\ell(\tau)$ can be compared with yield co-movement without numerically differentiating the empirical yield curve. Yet the price of improved numerical stability is that the integration in (12.4) can conceal features visible only in the instantaneous forward loading $B'(\tau)$. A yield-loading fit cannot by itself prove that instantaneous-forward volatility has a hump.

12.3 From a common factor to a Riccati maturity loading

Holding the parameters fixed, (12.6) gives the exact yield difference between two states:

$$y(\tau, r_2) - y(\tau, r_1) = \ell(\tau)(r_2 - r_1). \tag{12.8}$$

For maturities $\tau_1, \ldots, \tau_n$, write the loading vector as $\boldsymbol{\ell} = (\ell(\tau_1), \ldots, \ell(\tau_n))^\mathsf{T}$. Instantaneous yield innovations then have conditional covariance

$$\mathrm{Cov}_t(\,\mathrm{d}\boldsymbol{y}_t) = \beta r_t\,\boldsymbol{\ell}\boldsymbol{\ell}^\mathsf{T}\,\mathrm{d}t. \tag{12.9}$$

This matrix has rank one whenever $r_t > 0$ and $\boldsymbol{\ell} \neq 0$. In this exact model, all yield maturities respond to the same Brownian shock in fixed proportions. Empirical yields need not obey that exact restriction: additional risk factors, policy regimes, measurement errors and fitting procedures can generate extra eigenvalues.

There are at least three distinguishable statements: (i) a first principal component explains a high proportion of yield variation; (ii) its maturity profile resembles a Riccati loading $B(\tau)/\tau$; and (iii) a time-homogeneous one-state square-root model with the corresponding conditional state-dependent innovation variance and pricing relationships actually holds. The first is much weaker than the third.

A particularly transparent benchmark is

$$\ell_{\text{flat}}(\tau) \equiv 1. \tag{12.10}$$

This says that yield changes move equally across maturities. It has no fitted loading-shape parameters. If a flexible Riccati loading gives only a tiny improvement over (12.10), we should not interpret its fit as identification of the Riccati parameters or the sign of a.

12.4 What the fitted U.S. Treasury diagnostics actually show

The companion Version 16 paper reports an illustrative historical comparison using a monthly rolling-curve workbook derived from the Federal Reserve's nominal Treasury yield-curve estimates. It contains 538 month-end observations from September 1981 to June 2026, each formed by averaging the preceding 60 daily Svensson-fitted zero-coupon curves. The diagnostic uses demeaned, overlapping three-month changes across maturities from 1 to 15 years. These are *fitted* yields, not directly observed instantaneous forwards, independently measured bond-return risk premia or recovered stochastic volatilities. We reproduce the reported results here; no new estimation is undertaken for this chapter.

Table 12.1: Explained pooled variation in fitted zero-coupon yield changes, 1–15 years (three-month overlapping changes; percent). From the Version 16 companion paper.

Sample	Flat	$a > 0$ 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

The full-sample Riccati loading explains 91.19% of pooled squared fitted yield changes, but the flat-loading benchmark already explains 90.56%. The improvement is only 0.63 percentage points. The nonpositive-a branch gives the same displayed full-sample fit, while the unrestricted first principal component gives 91.24%. In other words, the data are strongly suggestive of a common maturity-movement component in the *fitted curves*, but they do not provide persuasive discrimination between the rival Riccati branches.

The 2010–2019 interval appears more favourable to positive a: the reported fit rises from 88.11% under the flat benchmark to 93.85% under the humped branch. That deserves attention, not immediate endorsement. At and near the zero lower bound, forward guidance and other monetary-policy measures can alter shorter-maturity yield variability, producing loading profiles that resemble the positive-a model without establishing that its risk-neutral diffusion is correct. This is a possible explanation, not a causal finding. The empirical maturity grid also starts at one year, so it does not observe the model's local slope $\ell'(0+) = a/2$.

What cannot be claimed. A 91.19% fit is not a validation of HJM pricing, exact rank-one covariance, physical stationarity or the volatility hump. Approximately 9% of the fitted-change variation remains outside the one-factor projection, and the exercise does not test whether that residual is merely measurement noise. The flat benchmark is also a limiting loading shape rather than a finite strictly positive-β Riccati parameter pair.

12.5 Why the stationary-mean proxy cannot reject monotonicity

Chapter 9 introduced the *optional* condition that the physical stationary expected forward curve be globally nondecreasing. Under that additional assumption, the humped model implies

$$\theta_P \le \theta_L = \frac{m}{\sqrt{a^2 + 2\beta}}, \qquad u = \frac{\theta_P}{\rho} \le \frac{1 - \sqrt{1 - 1/R}}{2}, \tag{12.11}$$

where $\rho = mB_\infty$ is the asymptotic forward rate and $R = 1 + a^2/(2\beta) > 1$ is the hump ratio. Here u is a dimensionless ratio, not an observed short-rate series.

The companion paper computes historical averages of the *fitted* one-year yield and fifteen-year instantaneous forward as illustrative short/long-tenor proxies:

Table 12.2: Descriptive short-/long-maturity proxy comparison. Fitted curve averages, 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

The full-sample ratio is 64.9%, exceeding one half. It would be a mistake to declare this a rejection of (12.11). The bound concerns θ_P/ρ, which neither sample column measures. Even in an ideal stationary population mean curve that is globally nondecreasing, its one-year average yield lies *above* the mean short rate θ_P, while its finite-maturity forward lies *below* the asymptotic rate ρ. For positive quantities, their ratio is therefore an *upper bound* on θ_P/ρ. A large upper bound cannot reject a requirement that the unobserved quantity itself be small. Moreover, pooled observations across changing regimes need not estimate a single invariant physical law.

The ratio is informative as an illustration of the measurement challenge, not as a test statistic. The 2020–2026 figure of 73.5% uses unrounded source averages; rounding the displayed numerator and denominator changes the ratio slightly.

12.6 Why the volatility hump is unusually hard to observe

Our theoretical result concerns the *instantaneous* conditional forward-volatility coefficient

$$q(\tau, r) = \sqrt{\beta r}\, B'(\tau). \tag{12.12}$$

A time series of changes in estimated instantaneous forward rates is not automatically a measurement of q. In particular, rolling-maturity and fixed-calendar-maturity dynamics have different drift transport terms; empirical finite differences also mix drift, changing volatility, measurement

error and possible regime shifts. Directly differentiating a Svensson curve twice—first to obtain forwards and then to estimate maturity-varying volatility from their changes—may amplify fitting artefacts.

Our Riccati loading provides an exact hump maturity and peak ratio on the $a > 0$ branch:

$$\tau_* = \frac{1}{\gamma} \log\left(\frac{\gamma + a}{\gamma - a}\right), \qquad R = 1 + \frac{a^2}{2\beta}, \qquad \gamma = \sqrt{a^2 + 2\beta}. \tag{12.13}$$

But the empirical question is whether the *observed instantaneous-forward covariance* is consistent with (12.12), not simply whether a smoothed volatility plot appears humped. A decisive test should use multiple observation intervals and curve-construction methods, control the maturity range actually supported by prices, and quantify uncertainty in the fitted hump's position and magnitude. Otherwise the plotted curvature may reflect the yield-fitting specification rather than the stochastic diffusion.

12.7 Expected excess returns offer another identification route

The preceding yield diagnostics depend almost entirely on the geometry of co-movement. Chapter 7 suggests a different source of information: instantaneous *conditional expected excess holding returns* of zero-coupon bonds. For the fixed-calendar-maturity bond with remaining tenor τ,

$$\mathrm{ERP}_\tau(r) = B(\tau)[b_Q(r) - b_P(r)]. \tag{12.14}$$

Two distinct empirical implications follow. First, whenever $B(\tau) > 0$,

$$\frac{\mathrm{ERP}_{\tau_1}(r)}{B(\tau_1)} = \frac{\mathrm{ERP}_{\tau_2}(r)}{B(\tau_2)}. \tag{12.15}$$

Second, after identifying the pricing parameters and the state-dependent expected excess return, the physical drift can be recovered algebraically:

$$b_P(r) = m + ar - \frac{\mathrm{ERP}_\tau(r)}{B(\tau)}. \tag{12.16}$$

The recovered b_P should be *independent of which bond maturity* is used. This is a sharp internal consistency requirement. Under zero-current physical stationarity, it further determines a candidate invariant density:

$$\pi_P(r) = \frac{C}{\beta r} \exp\left(\int^r \frac{2b_P(u)}{\beta u}\, \mathrm{d}u\right), \tag{12.17}$$

subject to normalisation, boundary classification, recurrence and a valid measure change.

There is, however, a crucial measurement obstacle. The instantaneous conditional expected excess return in (12.14) is *not* an observed realised excess return, an ex-post return average across all states or a conventional yield term premium. Estimating it requires forecasting the holding-period return conditional on the current state, adjusting for duration changes as the bond rolls down the curve, and accounting for sampling uncertainty. Finite-horizon expected excess returns need not obey the exact instantaneous formula without a corresponding holding-period adjustment.

Promising but untested. The cross-maturity proportionality of conditional instantaneous ERP offers a potential route to identifying physical drift. The historical yield-loading exercise above does not estimate ERP and therefore provides no evidence for or against (12.15).

12.8 A practical sequence of increasingly demanding tests

A useful future empirical programme would proceed in stages rather than fit all restrictions simultaneously. The stages below are methodological proposals, not results already obtained.

1. **Start with transparent price observations.** Construct consistent zero-coupon yield and price panels from underlying instruments where possible. Report interpolation, maturity coverage, transaction conventions and uncertainty, rather than regarding a fitted curve as direct observation of all tenors.

2. **Examine factor dimensionality.** Test the covariance rank of yield innovations within comparatively homogeneous regimes and across sampling frequencies. Distinguish an empirically dominant PC1 from the exact rank-one restriction of the model.

3. **Compare maturity loadings out of sample.** Fit the flat shape, both Riccati branches and unrestricted factor loadings on training periods; evaluate held-out periods, parameter stability and sensitivity to the short end. The positive branch must outperform a credible simpler alternative, not merely explain much of the total variation.

4. **Investigate state-dependent variance.** Conditional on an independently observed short-rate state, test whether yield-innovation covariance scales as $\beta r_t \boldsymbol{\ell}\boldsymbol{\ell}^{\mathsf{T}}$. This is stronger than matching the maturity shape of the loading.

5. **Estimate physical drift and expected returns independently.** Recover short-rate conditional forecasts and bond expected excess returns with explicit horizon conventions; test the maturity-invariance in (12.15) and consistency of the reconstructed drift in (12.16).

6. **Only then examine stationary geometry.** Within defensibly stable physical regimes, estimate or bound θ_P and ρ; test the optional global monotonicity selection rather than treating it as an arbitrage condition.

The order matters. Later tests inherit the uncertainty of earlier construction and identification choices. A highly constrained joint calibration that fits prices and returns may still be uninformative if key restrictions are effectively imposed rather than tested.

12.9 Regime dependence, falsification and the limits of one state

An interest-rate state can be sensitive to central-bank policy, changing inflation expectations, liquidity conditions and the composition of the Treasury market. A full historical series may mix different physical regimes. If the short-rate drift, market price of risk or effective number of factors changes across those regimes, it is inappropriate to interpret pooled observations as draws from one time-homogeneous stationary law. This matters especially for the low-rate interval, when the one-year to fifteen-year shape may be heavily influenced by policy implementation and zero-lower-bound effects.

The mathematical results in Chapters 3–11 are conditional. A failure of exact rank-one covariance in observable yield innovations would challenge the empirical adequacy of an exact one-state model for those data. It would not refute Itô calculus, HJM no-arbitrage or the algebra of the affine pricing PDE. Likewise, failure of the global mean bound would challenge the joint package of the selected pricing branch, faithful measurement and *optional* monotonic equilibrium geometry; it would not disprove physical stationarity in general.

The converse matters too. Even if a fitted maturity loading resembles $B(\tau)/\tau$, that does not prove the model's full P–Q compatibility, a particular physical invariant density or an admissible

market-price-of-risk change of measure. Model validation requires evidence that is not contained in the fitted loading shape alone.

Scope of the historical evidence. The numerical tables above reproduce the companion Version 16 paper's diagnostic calculations. The present chapter does not rerun the workbook, supply confidence intervals, estimate state-conditional expected bond returns, or constitute a formal hypothesis test. Its principal contribution is to clarify what the existing numbers can and cannot support.

12.10 What this chapter establishes, and what remains free

The illustrative Treasury diagnostics show that a Riccati yield loading can summarise much of the movement of fitted yield curves. They do *not* identify the positive-a risk-neutral branch: the full-sample improvement over a flat loading is small and both Riccati branches fit almost identically. The short-/long-tenor proxy ratio cannot reject the optional monotonic mean restriction because it does not measure the required physical stationary mean and asymptotic forward rate. These negative identification results are valuable: they prevent empirical description from being mistaken for mathematical verification.

Our framework nevertheless supplies testable restrictions. Exact one-state covariance rank, the nonlinear Riccati maturity loading, square-root state-dependent innovation variance and cross-maturity proportionality of instantaneous expected excess bond returns form an increasingly demanding empirical hierarchy. Independent expected-return and physical forecast information might eventually narrow the class of admissible physical drifts. Until such evidence exists, the mathematics establishes *possible coexistence*, not an empirical claim that observed Treasury markets follow this particular one-state diffusion.

The final chapter returns to the original motivation: what has been unified, which assumptions remain conditional, and where the most productive further research may lie.

Chapter 13

Conclusions and Open Questions

The preceding chapters develop a conditional compatibility framework for a one-state interest-rate term structure. Exact forecast closure selects affinity of the state-to-forward-curve map under appropriate nondegeneracy and regularity conditions. Arbitrage-free bond pricing then constrains the risk-neutral state dynamics; a positive square-root diffusion and an interior instantaneous-forward volatility hump select a non-restoring risk-neutral affine drift. Neither this pricing geometry nor physical stationarity determines a unique physical drift. The market price of risk reconciles the two, while expected excess bond returns provide a direct route to identifying the otherwise free physical drift if those returns can be measured independently. Stationary averaging and the long-forward measure yield a further, explicitly optional, mean restriction when the equilibrium forward curve is required to be globally nondecreasing. We close by distinguishing mathematical coexistence from empirical validation and by identifying questions that could sharpen or challenge the framework.

13.1 Returning to the original question

A yield curve can be described in several ways. It is a collection of market prices; a maturity-indexed field of instantaneous forward rates; a map from a stochastic state to future possible curves; an object whose evolution must satisfy no-arbitrage; and, when the physical state is stationary, an object with a stationary expected shape. These descriptions are related, but they are not interchangeable. In particular, a physical forecast is not a risk-neutral price, a stationary average is not a pathwise invariant curve, and a volatility loading is not itself a measured volatility without a stochastic state scale.

The purpose of this monograph has been to ask a deliberately restricted question: *can these descriptions coexist inside a mathematically precise one-state model, and which restrictions does their coexistence impose?* We did not assume that all term structures should have one stochastic factor. Nor did we claim that empirical interest-rate markets follow the selected model. The task has instead been to expose the logical relationships among hypotheses which are often developed separately.

The result is a hierarchy, not a universal dynamics. Its strength is that each additional requirement changes what may be concluded, while allowing us to say where the classification stops.

13.2 The hierarchy in one view

A concise description of the model begins under the physical probability measure P:

$$\mathrm{d}r_t = b_P(r_t)\,\mathrm{d}t + \sqrt{\beta r_t}\,\mathrm{d}W_t^P, \qquad r_t > 0, \quad \beta > 0. \tag{13.1}$$

At the outset, the function b_P is not required to be affine. The forward curve is written $x(\tau, r)$, with $x(0, r) = r$. We impose the *forecast-closure selection*

$$\mathbb{E}_r^P[x(\tau, r_h)] = x\Big(\tau, \mathbb{E}_r^P[r_h]\Big). \tag{13.2}$$

Under the regularity conditions developed in Chapter 3, nondegenerate short-rate variance forces the second state derivative of x to vanish. Consequently, the curve is affine in the state,

$$x(\tau, r) = c(\tau) + d(\tau)r. \tag{13.3}$$

Equation (13.2) is *not* a universal identity for conditional expectations. It is an economically motivated modelling selection: the expected state is taken to be sufficient to reproduce the expected curve.

When an exponential-affine bond price $P(\tau, r) = \exp[-A(\tau) - B(\tau)r]$ is required to satisfy risk-neutral pricing for a sufficiently rich range of maturities, the risk-neutral drift and variance must be affine functions of the state. Selecting a vanishing diffusion variance at the origin gives the square-root branch, and the pricing equations become

$$\mathrm{d}r_t = (m + ar_t)\,\mathrm{d}t + \sqrt{\beta r_t}\;\mathrm{d}W_t^Q, \tag{13.4}$$

$$A'(\tau) = mB(\tau), \qquad B'(\tau) = 1 + aB(\tau) - \frac{\beta}{2}B(\tau)^2, \qquad A(0) = B(0) = 0. \tag{13.5}$$

These equations also follow from HJM's restriction when the rolling maturity transport is handled correctly. The two routes must agree because they describe the same traded-bond prices; they are not independent no-arbitrage laws.

The logical point. Forecast closure, affine pricing, zero variance at the short-rate boundary, and global monotonicity of the physical equilibrium forward curve are *different restrictions*. Their consequences should not be compressed into one statement such as "one state implies the Riccati model."

13.3 The hump and the apparent conflict with equilibrium

The instantaneous-forward curve and its diffusion loading under (13.4) are

$$x(\tau, r) = mB(\tau) + B'(\tau)r, \qquad q(\tau, r) = \sqrt{\beta r}\,B'(\tau). \tag{13.6}$$

Differentiate the Riccati equation:

$$B''(\tau) = [a - \beta B(\tau)]B'(\tau). \tag{13.7}$$

For the positive loading solution, $B(0) = 0$, $B' > 0$, and B increases toward its finite positive root. If $a \le 0$, B' decreases from the short boundary. If $a > 0$, B' initially rises, reaches a unique interior maximum at $B(\tau_*) = a/\beta$, and then declines. Thus, *within this selected one-state square-root affine class*, an interior instantaneous-forward-volatility hump is equivalent to $a > 0$.

Write

$$\gamma = \sqrt{a^2 + 2\beta}, \qquad B_\infty = \frac{a + \gamma}{\beta}. \tag{13.8}$$

The hump's height relative to the short-end loading and its location are

$$R = 1 + \frac{a^2}{2\beta}, \qquad \tau_* = \frac{1}{\gamma}\log\left(\frac{\gamma + a}{\gamma - a}\right) \quad (a > 0). \tag{13.9}$$

These formulas are consequences of a familiar Riccati equation. Their significance is in what they select within the compatibility hierarchy, not in a claim that the Riccati equation itself is new.

The selected risk-neutral drift $m + ar$ is non-restoring. That seems at odds with a stationary short rate only if the physical and pricing measures are confused. They need not have the same drift. Indeed, their difference is precisely the instantaneous risk-price adjustment:

$$b_Q(r) - b_P(r) = \sigma(r)\lambda(r), \qquad \sigma(r) = \sqrt{\beta r}. \tag{13.10}$$

The risk-neutral process need not possess an invariant distribution for the physical process to possess one. A suitable physical drift may be restoring, subject to positivity, recurrence, and validity of the change of measure.

13.4 Risk compensation provides the compatibility mechanism

Consider a zero-coupon bond with current remaining maturity $\tau > 0$. Its *instantaneous conditional expected excess holding-return rate* under P is

$$\mathrm{ERP}_\tau(r) = B(\tau)\sigma(r)\lambda(r). \tag{13.11}$$

It follows immediately from (13.10) that

$$\boxed{\mathrm{ERP}_\tau(r) + B(\tau)b_P(r) = B(\tau)(m + ar).} \tag{13.12}$$

Here we meet the central theme of the monograph in a particularly economical form. The pricing model fixes the combination of the physical forecast contribution and the risk-compensation contribution, but does not fix either separately. If the physical drift has nonlinear state dependence, the risk-premium function must offset that nonlinearity in exactly the way required to preserve the affine pricing drift.

There is an exact maturity restriction as well. For any $\tau_1, \tau_2 > 0$ with nonzero loadings,

$$\frac{\mathrm{ERP}_{\tau_1}(r)}{B(\tau_1)} = \frac{\mathrm{ERP}_{\tau_2}(r)}{B(\tau_2)} = m + ar - b_P(r). \tag{13.13}$$

Thus the state dependence of instantaneous expected excess returns is shared across maturities, up to the Riccati loading. This is a strong restriction of the one-factor model, though it is already implicit in standard affine bond-return formulas. It should not be confused with an identical property of *realised* or finite-horizon bond returns.

If an independent estimate of the conditional expected excess return were available, (13.12) could be inverted:

$$\boxed{b_P(r) = m + ar - \frac{\mathrm{ERP}_\tau(r)}{B(\tau)}.} \tag{13.14}$$

The maturity used to recover the physical drift should not affect the answer. This is both an identification opportunity and a falsifiable implication. In practice, conditional instantaneous ERP is latent and its estimation is difficult; the algebra does not solve the measurement problem.

13.5 Physical equilibrium is flexible, but not unconstrained

Suppose the physical diffusion admits an invariant density π_P with no stationary probability flux at the boundaries. The stationary Fokker–Planck equation gives

$$b_P(r) = \frac{\beta}{2}\left[1 + r\,\partial_r \log \pi_P(r)\right]. \tag{13.15}$$

Conversely, a candidate density can be reconstructed from an admissible physical drift:

$$\pi_P(r) = \frac{C}{\beta r} \exp\left[\int^r \frac{2b_P(u)}{\beta u}\, \mathrm{d}u\right]. \tag{13.16}$$

Normalisability alone is not an adequate substitute for boundary classification, non-explosion, recurrence and the measure-change hypotheses. Under those conditions, however, neither the pricing Riccati equation nor the volatility hump singles out the Gamma family. The restoring CIR case is one example; nonlinear restoring drifts can produce non-Gamma stationary laws compatible with the same pricing geometry.

Let the stationary physical mean be $\theta_P = \mathbb{E}_{\pi_P}[r]$. Since the forward curve is affine,

$$x_e^P(\tau) = \mathbb{E}_{\pi_P}[x(\tau, r)] = mB(\tau) + \theta_P B'(\tau). \tag{13.17}$$

Thus the *entire shape* of the physical expected equilibrium forward curve depends on the stationary distribution only through its first moment. Two distinct stationary distributions with the same mean produce the same x_e^P. This is an exact non-identification result, not an approximation.

Under stationarity $\mathbb{E}_{\pi_P}[b_P] = 0$. Combining this with (13.12) yields

$$\mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau] = B(\tau)(m + a\theta_P), \tag{13.18}$$

$$x_e^P(\tau) - \theta_P = \mathbb{E}_{\pi_P}[\mathrm{ERP}_\tau] - \frac{\beta\theta_P}{2} B(\tau)^2. \tag{13.19}$$

The equilibrium forward spread therefore combines mean bond risk compensation and a volatility correction. In particular, the short-boundary slope obeys

$$\boxed{\partial_\tau x_e^P(0) = m + a\theta_P = \mathbb{E}_{\pi_P}[\sigma(r)\lambda(r)].} \tag{13.20}$$

The equality concerns the *physical stationary average* of a risk-neutral drift; it does not assert that Q has a stationary distribution. Nor should it be read as a finite-horizon equality between physical and risk-neutral expected short-rate levels.

13.6 The optional long-forward selection

The long-forward pricing eigenfunction for the selected affine model is proportional to $\exp[-B_\infty r]$. Its associated measure, denoted L, has restoring drift

$$b_L(r) = m - \gamma r, \qquad \theta_L = \frac{m}{\gamma}. \tag{13.21}$$

This provides a useful comparison mean, but does not automatically constrain the mean of the physical invariant law. An additional geometric requirement is needed.

Differentiate (13.17) and use the Riccati equation:

$$\partial_\tau x_e^P(\tau) = B'(\tau)\big[m + a\theta_P - \beta\theta_P B(\tau)\big]. \tag{13.22}$$

For $a > 0$, $B' > 0$ and B increases to B_∞. Therefore, the *optional* requirement that the stationary expected forward curve be nondecreasing at every maturity is equivalent to

$$\boxed{\theta_P \le \theta_L = \frac{m}{\sqrt{a^2 + 2\beta}}.} \tag{13.23}$$

The consequence is a bound on a mean, not the selection of a unique stationary density. It should not be described as a necessary property of physical equilibrium or arbitrage-free pricing.

For clarity, if $\rho = mB_\infty$ denotes the asymptotic forward rate and $u = \theta_P/\rho$, the humped branch's monotonicity restriction can be written

$$0 < u < \frac{1}{2}, \qquad 1 < R \leq \frac{1}{4u(1-u)}. \tag{13.24}$$

This is an observable-looking inequality, but meaningful application still requires reliable identification of the physical stationary mean, the asymptotic forward rate and the hump ratio. The historical proxy comparisons in Chapter 12 do not establish those measurements.

13.7 What the framework unifies—and what it does not

The relationships above use several well-established mathematical theories: the Markov generator and conditional forecasts, the affine bond-pricing equation, the HJM drift restriction, Girsanov measure changes, bond risk-premium identities, stationary diffusion theory and long-forward pricing eigenfunctions. Their provenance matters. This monograph does not claim to have discovered these ingredients independently of the literature.

Its more defensible contribution is a *joint conditional classification*: a precise account of which ingredients can hold together, which apparent conflicts are resolved by changes of measure, which extra assumptions select narrower subclasses, and which physical quantities remain free. The structure is restrictive about pricing geometry and maturity loadings, but permissive about the physical stationary law. That permissiveness is not a failure of the framework: it is an identification result.

Requirement or observation	What it establishes within the selected class
Exact forecast closure	Affine state-to-curve representation under nondegeneracy and regularity.
Arbitrage-free affine pricing	Affine risk-neutral generator and Riccati maturity loading.
Vanishing variance at $r = 0$	Square-root diffusion branch.
Interior forward-volatility hump	Positive risk-neutral state drift coefficient a.
Physical stationarity	A physical invariant law, not a unique b_P or density.
Independent state-conditional ERP	A potential identification route for b_P and π_P.
Global monotonicity of x_e^P	Optional physical-mean bound $\theta_P \leq \theta_L$.

The scope of the conclusions. The mathematical results are conditional on the assumed one-state Markov structure, exact forecast closure, suitable diffusion regularity, an admissible change of measure, and the selected affine pricing class. The empirical comparisons are preliminary and do not establish that actual interest-rate markets satisfy all these requirements.

13.8 Open questions and a disciplined next stage

The natural research programme follows from the distinctions the hierarchy has made visible.

Identification from expected returns. Can conditional instantaneous excess bond-return rates be inferred robustly enough to test (13.13)? That test would require defensible holding-return horizon conventions, high-quality bond data, and uncertainty quantification. Because finite-horizon returns mix risk compensation with subsequent state evolution and remaining-maturity changes, instantaneous identities must not be applied to them without adjustment.

Independent identification of physical dynamics. Can physical short-rate forecasts, estimated from a sufficiently stable regime, recover a drift function that agrees with the one implied by bond excess returns? Agreement across maturities and forecast horizons would be more informative than fitting the yield curve alone.

The volatility hump and model selection. Can the positive-a branch be distinguished out of sample from the nonpositive branch, a flat loading and multifactor competitors? Chapter 12 showed why fitted yields alone have not yet done so. A sharp test should target state-dependent covariance and short-maturity volatility, not merely the proportion of variance explained by a first principal component.

Stationary-law admissibility. Which classes of nonlinear restoring physical drift admit globally valid positive diffusions, normalisable zero-current invariant laws and valid equivalent changes of measure to the non-restoring pricing branch? The local algebra $b_Q - b_P = \sigma\lambda$ does not by itself settle infinite-horizon equivalence or recurrence.

Regimes, interventions and more than one state. A single time-homogeneous physical drift may be inadequate across policy regimes, lower-bound episodes and large structural transitions. If extra states are introduced, exact forecast closure, covariance rank and cross-maturity ERP factorisation must be reconsidered. Such extensions should be driven by a clearly identified empirical failure, rather than introduced merely to improve fit.

The optional equilibrium geometry. Is there a compelling economic reason for the *stationary expected* forward curve to be globally nondecreasing? If not, the bound (13.23) should remain a useful conditional selection, not an assumption elevated into a law. Better measurements of the physical stationary mean and long-forward quantities would help determine whether the selection is empirically informative.

None of these questions requires us to reverse the mathematical conclusions already established. They concern either model admissibility beyond formal local calculations, the identification of currently latent quantities, or the suitability of the chosen class for observed rates.

13.9 Closing perspective

The strongest result of this investigation may be a distinction rather than a single formula. A model can impose sharp restrictions on the geometry of arbitrage-free prices without selecting a unique physical probability distribution. Conversely, physical equilibrium can coexist with a pricing measure whose state drift is non-restoring. The market price of risk explains the difference, and expected excess returns express it in the language of actual financial payoffs.

At the level of mathematics, the programme has therefore replaced a search for one universal dynamic law with a more tractable and testable question: *which descriptions of the interest-rate system can be true simultaneously, and what additional information would tell them apart?* This is the sense in which the framework unifies established results. It does not erase the distinctions between physical forecasting, risk-neutral pricing, stationary equilibrium and long-horizon valuation. It makes those distinctions explicit, and then derives their points of compatibility.

Final principle. Compatibility specifies how descriptions must fit together. Identification specifies what can be learned from observations. The distinction is not a weakness to be concealed; it is what allows the model's genuine mathematical restrictions to be separated from assumptions of convenience and claims that must still be tested.

References

Bibliography

[1] Hansen, L. P., and Scheinkman, J. A. (2009). Long-term risk: An operator approach. *Econometrica*, 77, 177–234.

[2] Qin, L., and Linetsky, V. (2016). Positive eigenfunctions of Markovian pricing operators: Hansen–Scheinkman factorization, Ross recovery, and long-term pricing. *Operations Research*, 64, 99–117.

[3] Duffie, D., and Kan, R. (1996). A yield-factor model of interest rates. *Mathematical Finance*, 6, 379–406.

[4] Duffee, G. R. (2002). Term premia and interest rate forecasts in affine models. *Journal of Finance*, 57, 405–443.

[5] 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, 123–170.

[6] Board of Governors of the Federal Reserve System. Nominal yield curve data and documentation. `https://www.federalreserve.gov/data/nominal-yield-curve.htm`.

[7] Svensson, L. E. O. (1994). Estimating and interpreting forward interest rates: Sweden 1992–1994. NBER Working Paper 4871.

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

[9] Cochrane, J. H., and Piazzesi, M. (2005). Bond risk premia. *American Economic Review*, 95, 138–160.

[10] Heath, D., Jarrow, R., and Morton, A. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 60, 77–105.

[11] Piazzesi, M. (2010). Affine term structure models. In *Handbook of Financial Econometrics*, Vol. 1, pp. 691–766.