IRFR Monograph

A Boundary-Field Theory of Interest Rates

Contents

1 Introduction and Programme **1**
1.1 Motivation . . . 1
1.2 The central principle . . . 2
1.3 The IRFR primitive . . . 3
1.4 HJM: representation versus explanation . . . 4
1.5 Boundary laws and monetary policy . . . 5
1.6 Comparison with existing literature . . . 6
1.7 Claims and limitations . . . 6
1.8 Roadmap . . . 7

2 The IRFR Primitive and the Field Representation **9**
2.1 Introduction . . . 9
2.2 From calendar maturity to rolling maturity . . . 10
2.3 The transport identity . . . 11
2.4 The forward curve as a maturity-domain field . . . 11
2.5 Local rate exposure and the economic meaning of the IRFR coordinate . . . 12
2.6 Transport toward the boundary . . . 13
2.7 The economic meaning of the short-end boundary . . . 14
2.8 Relation to the conventional forward curve . . . 14
2.9 Why the primitive matters . . . 15
2.10 Summary . . . 15
2.11 Transition to the axioms . . . 15

3 The IRFR Axioms **17**
3.1 Introduction . . . 17
3.2 Why axioms? . . . 18
3.3 Axiom 1: stochastic evolution . . . 18
3.4 Interpretation of Axiom 1 . . . 19
3.5 Axiom 2: rolling-maturity invariance . . . 19
3.6 Interpretation of Axiom 2 . . . 20
3.7 Axiom 3: equilibrium moment transport . . . 20
3.8 Running example: a sloped equilibrium curve . . . 21
3.9 Interpretation of Axiom 3 . . . 22
3.10 Axiom 4: boundary closure . . . 22
3.11 Interpretation of Axiom 4 . . . 22
3.12 Interior law versus boundary law . . . 23

3.13 Summary and transition . . . 23

4 The Kolmogorov Forward Equation and Affine Closure **25**
4.1 Introduction . . . 25
4.2 From axioms to a density equation . . . 26
4.3 Moment equations . . . 26
4.4 The transport correction . . . 27
4.5 Affine closure of the drift operator . . . 28
4.6 Interpretation of the affine drift . . . 29
4.7 Sign conventions and mean reversion . . . 30
4.8 Variance and common maturity loading . . . 31
4.9 Moment closure and admissibility . . . 31
4.10 Relation to HJM . . . 32
4.11 Summary . . . 32
4.12 Transition to admissible modes . . . 33

5 Admissible Maturity Modes and the r-Branch **34**
5.1 Introduction . . . 34
5.2 From affine closure to admissibility . . . 35
5.3 Coefficient matching . . . 36
5.4 The characteristic equation . . . 36
5.5 The r-parameter . . . 37
5.6 Distinct-root branch . . . 38
5.7 Repeated-root branch . . . 38
5.8 The r = 1 branch . . . 39
5.9 Relation to humped volatility . . . 40
5.10 Interior modes and boundary selection . . . 40
5.11 Relation to Ritchken–Chuang . . . 41
5.12 Summary . . . 41
5.13 Transition to boundary closure . . . 42

6 Boundary Closure and Economic Regimes **43**
6.1 Introduction . . . 43
6.2 Why boundary closure is necessary . . . 44
6.3 Interior admissibility versus boundary selection . . . 44
6.4 The short end as an economic boundary . . . 45
6.5 Three canonical boundary regimes . . . 45
6.6 Neumann/free closure . . . 46
6.7 Dirichlet closure . . . 47
6.8 The Dirichlet repeated-root branch . . . 48
6.9 Markovianity and representation . . . 48
6.10 Robin closure . . . 49
6.11 Boundary law as policy transmission . . . 50
6.12 Relation to the r-branch . . . 50
6.13 Summary . . . 51
6.14 Transition to detailed boundary cases . . . 51

7 Neumann/free Boundary Closure and the Free Benchmark **52**

7.1 Introduction . . . 52
7.2 What "Neumann/free" means . . . 52
7.3 The short end under Neumann/free closure . . . 53
7.4 Neumann/free closure and probability flux . . . 54
7.5 Neumann/free closure and the free stochastic mode . . . 54
7.6 Benchmark interpretation . . . 55
7.7 Neumann/free closure versus Dirichlet selection . . . 55
7.8 Neumann/free closure and the benchmark economy . . . 56
7.9 Consequences for volatility shape . . . 56
7.10 Relation to monetary policy . . . 57
7.11 Summary . . . 57
7.12 Transition to Dirichlet closure . . . 57

8 Dirichlet Boundary Closure and the r = 1 Humped Branch 58
8.1 Introduction . . . 58
8.2 Dirichlet closure as fixed-boundary selection . . . 59
8.3 Economic interpretation: the pinned short end . . . 59
8.4 Action on the repeated-root mode space . . . 59
8.5 The selection calculation . . . 60
8.6 Shape of the selected loading . . . 61
8.7 Positivity and the lower admissible boundary . . . 61
8.8 Comparison with Neumann/free closure . . . 62
8.9 Markovianity of the Dirichlet humped branch . . . 63
8.10 Relation to Ritchken–Chuang-type volatility . . . 63
8.11 Why the Dirichlet case is extreme . . . 64
8.12 Boundary law and policy transmission . . . 64
8.13 Summary . . . 64
8.14 Transition to Robin closure . . . 65

9 Robin Boundary Closure and Partial Policy Anchoring 66
9.1 Introduction . . . 66
9.2 Robin closure as a mixed boundary condition . . . 67
9.3 Economic interpretation: partial anchoring . . . 67
9.4 Action on the repeated-root mode space . . . 68
9.5 Limiting cases and reference regimes . . . 69
9.6 The anchoring parameter . . . 69
9.7 Shape of the Robin-selected loading . . . 70
9.8 Positivity and admissibility under Robin closure . . . 71
9.9 Robin closure and monetary-policy regimes . . . 71
9.10 Comparison with Neumann/free and Dirichlet regimes . . . 72
9.11 Relation to humped volatility . . . 72
9.12 Boundary law as a regime parameter . . . 73
9.13 Summary . . . 73
9.14 Transition . . . 74

10 Boundary-Selected Loadings and Observable Volatility 75
10.1 Introduction . . . 75
10.2 From diffusion scale to volatility loading . . . 76

10.3 Boundary-selected loading families . . . 76
10.4 The Neumann/free exponential benchmark . . . 77
10.5 The Dirichlet humped loading . . . 78
10.6 The Robin mixed loading . . . 79
10.7 Shape diagnostics of the Robin family . . . 80
10.8 Sign compatibility and lower-boundary interpretation . . . 81
10.9 Comparison with Ritchken–Chuang-type forms . . . 81
10.10 Quadratic models and boundary-selected loadings . . . 82
10.11 Representation in HJM coordinates . . . 82
10.12 Empirical interpretation of boundary regimes . . . 83
10.13 The role of the r=1 branch . . . 83
10.14 Summary . . . 84

11 Stochastic Long-Maturity Boundary States **86**
11.1 Motivation . . . 86
11.2 Boundary decomposition . . . 87
11.3 Single-risk-direction admissibility . . . 88
11.4 Asymptotic boundary dynamics . . . 88
11.5 Generator corroboration . . . 89
11.6 Expected boundary reversion . . . 89
11.7 Long-end variance floor . . . 90
11.8 Curve representation . . . 90
11.9 Branch specialisations . . . 91
11.10 Relation to risk-neutral pricing . . . 92
11.11 Summary . . . 92

12 Bond Pricing and the Risk-Neutral IRFR Measure **93**
12.1 Introduction . . . 93
12.2 The money-market account at an initial date . . . 93
12.3 Generalisation to an arbitrary time . . . 94
12.4 The IRFR strip and zero-coupon bond prices . . . 95
12.5 Stochastic consistency of rolling replication . . . 95
12.6 Physical and pricing dynamics . . . 96
12.7 Relation to the HJM drift condition . . . 97
12.8 Boundary interpretation of the short rate . . . 99
12.9 Derivative-pricing extension . . . 99
12.10 What this construction does and does not prove . . . 100

13 Calibration, Identification, and Empirical Discipline **101**
13.1 Introduction . . . 101
13.2 Principle 1: estimate the loading, not only the volatility level . . . 102
13.3 Principle 2: separate stochastic dimension from maturity geometry . . . 103
13.4 Principle 3: identify the decay scale . . . 104
13.5 Principle 4: estimate boundary anchoring from short-end behaviour . . . 105
13.6 Principle 5: do not infer r = 1 from hump resemblance alone . . . 106
13.7 Principle 6: distinguish scale from shape . . . 106
13.8 Principle 7: use principal components carefully . . . 107
13.9 Principle 8: treat boundary regimes as estimable hypotheses . . . 108

13.10 Principle 9: check positivity and admissibility separately . . . 109
13.11 Principle 10: empirical fit is evidence, not proof . . . 110
13.12 Minimal one-factor IRFR calibration . . . 110
13.13 A practical calibration sequence . . . 111
13.14 Summary . . . 111
13.15 Transition . . . 112

14 Empirical Evidence for the Completed Boundary Geometry 113
14.1 Introduction . . . 113
14.2 Geometric Invariance and Dynamic Risk Weighting . . . 114
14.3 From the Completed Jacobian to an Empirical State . . . 115
14.4 Dataset and Empirical Construction . . . 116
14.5 Boundary-Risk Decomposition . . . 116
14.6 Market Price of Risk and MSNR . . . 117
14.7 Empirical Results . . . 117
14.7.1 Branch selection . . . 117
14.7.2 One Brownian versus two boundary-risk directions . . . 118
14.7.3 Spatial decay stability . . . 118
14.7.4 Fixed-a^2 volatility geometry . . . 118
14.7.5 MSNR and the risk-neutral bridge . . . 118
14.8 Interpretation . . . 118
14.9 Limitations . . . 119
14.10 Conclusion . . . 119

15 Scope, Limitations, and Further Directions 120
15.1 Introduction . . . 120
15.2 What has been established . . . 121
15.3 What has been supported, but remains conditional . . . 122
15.4 What has not been established . . . 123
15.5 Status of the r = 1 branch . . . 124
15.6 Status of the boundary-law interpretation . . . 124
15.7 Minimality and modelling cost . . . 125
15.8 Further directions . . . 125
15.9 Why the minimal model still matters . . . 126
15.10 Summary . . . 126
15.11 Transition to the conclusion . . . 127

16 Conclusion: A Unified Boundary-Field Theory of Interest Rates 128
16.1 Introduction . . . 128
16.2 The problem revisited . . . 129
16.3 The IRFR primitive . . . 129
16.4 The axioms . . . 130
16.5 The affine drift operator . . . 130
16.6 Admissible maturity modes . . . 131
16.7 Boundary selection . . . 132
16.8 Observable volatility . . . 132
16.9 What has been claimed . . . 133
16.10 What has not been claimed . . . 134

16.11 Why the framework matters . . . 134
16.12 Why the Theory Is Unified . . . 135
16.13 Open problems . . . 136
16.14 Final statement . . . 136

A Notation and Conventions **138**
A.1 Purpose of this appendix . . . 138
A.2 Time and maturity variables . . . 138
A.3 Forward rate and IRFR primitive . . . 138
A.4 Transport identity . . . 139
A.5 IRFR stochastic convention . . . 139
A.6 Kolmogorov forward equation convention . . . 139
A.7 Affine drift operator . . . 139
A.8 Diffusion loading and covariance rate . . . 140
A.9 Admissible maturity modes . . . 140
A.10 The r-branch . . . 140
A.11 Boundary regimes . . . 140
A.12 Lower-boundary convention . . . 141

B Algebra of Affine Closure, the r-Branch, and Boundary-Selected Loadings **142**
B.1 Purpose of this appendix . . . 142
B.2 Stochastic convention and KFE . . . 142
B.3 First moment . . . 142
B.4 Affine closure . . . 143
B.5 Diffusion loading . . . 143
B.6 The r-branch . . . 143
B.7 Distinct-root Dirichlet loading . . . 144
B.8 The repeated-root Dirichlet limit . . . 144
B.9 Dirichlet hump and variance rate . . . 144
B.10 Shifted inverse-gamma density . . . 145
B.11 Feller inaccessibility . . . 145
B.12 Neumann/free benchmark . . . 145
B.13 Robin loading and anchoring parameter . . . 146
B.14 Robin sign compatibility . . . 146
B.15 Summary . . . 146

References **147**

Chapter 1

Introduction and Programme

1.1 Motivation

Interest-rate modelling has developed through several powerful but partially distinct traditions. Short-rate models provide intuitive state variables and tractable dynamics [15, 3, 7]. Heath–Jarrow–Morton (HJM)-type models provide a broad arbitrage-free representation of forward-rate evolution [6]. Affine term-structure models impose analytical structure and calibration discipline [4]. Market models focus directly on traded rates and market observables [2].

Each of these traditions has been extraordinarily productive, yet a conceptual separation remains. In many models, the dynamic consistency of the term structure is handled in one part of the framework, while the shape of the curve, the maturity profile of volatility, and the effect of monetary policy are introduced elsewhere through additional parametrisation. Humped volatility structures, curve-shape restrictions, maturity loadings, and policy effects are often specified because they are useful or empirically plausible, rather than derived from a common structural mechanism.

The empirical importance of humped volatility makes this separation especially visible. Ritchken–Chuang-type maturity profiles are useful precisely because they capture an observed concentration of volatility away from the shortest and longest maturities [13, 14, 9, 10, 11]. Quadratic term-structure models and other nonlinear extensions provide one powerful response to this problem by enriching the state representation [1, 8].

The question pursued here is different: can some of the same empirically recognisable maturity structure be explained before adding richer polynomial state machinery, by first exploiting rolling-maturity geometry and the boundary at the short end?

The purpose of this monograph is to investigate whether more of this structure can be generated from fewer principles, while preserving a clear account of the modelling cost. The aim is not merely parsimony, but internal consistency: curve shape, volatility structure, maturity loading, and policy anchoring should be linked through the same governing field principles, so that every modelling restriction has an identifiable location in the theory.

The central object is the instantaneous rolling forward rate. Instead of taking the forward rate at calendar time t for fixed maturity T as the primitive object, the IRFR formulation works in rolling-maturity coordinates. Writing

$$x(t,\tau) = F(t,t+\tau), \tau = T - t,$$

the forward curve becomes a stochastic field over remaining maturity τ. Calendar time evolves the field, while the maturity coordinate rolls deterministically toward the short end. The point $\tau = 0$ is therefore not merely a limiting maturity. It is the boundary of the term-structure domain.

This change of primitive is not just a change of notation. It changes the geometry of the modelling problem. The curve is no longer viewed primarily as a family of unrelated forward-rate processes indexed by maturity. It is viewed as a maturity-domain field transported through time. Its admissible behaviour is constrained by stochastic evolution, rolling-maturity consistency, equilibrium moment transport, and boundary closure at the short end.

This is the sense in which the present work develops a field theory of interest rates. The phrase is not intended as a metaphor imported from physics for decorative effect. It refers to the treatment of the term structure as a stochastic object defined over a continuous maturity domain, with an interior law, admissible modes, and boundary conditions. The field is the forward-rate curve in rolling-maturity coordinates. The boundary is the short end. The boundary law is the rule that closes the system.

The guiding claim of the monograph is that a substantial part of term-structure behaviour can be reorganised around this structure. Curve shape, maturity loading, volatility localisation, and monetary-policy anchoring need not be introduced as separate modelling interventions. They can be interpreted as consequences of the same rolling-field geometry and its boundary laws.

1.2 The central principle

The organising principle of the monograph is:

Dynamics are determined by invariance; observable shapes are selected by boundary laws.

This sentence is deliberately schematic, so as to capture the architecture of the theory.

The first part of the statement concerns the interior of the maturity domain. Rolling-maturity invariance and equilibrium moment transport restrict the admissible stochastic dynamics of the IRFR field. The Kolmogorov forward equation and affine closure then force a particular relation between drift, diffusion, and maturity loading. The result is not an arbitrary volatility specification but a structured family of admissible maturity modes.

The second part of the statement concerns the boundary at $\tau = 0$. The interior dynamics determine what the field is allowed to do, but they do not by themselves determine which admissible mode is realised. That selection is made by the boundary law. A Neumann/free[1] boundary, a fixed boundary, and a partially anchored boundary select different maturity structures, and therefore different observable curve and volatility shapes.

This separation is central. The theory distinguishes three objects that are often conflated in term-structure modelling:

1. the stochastic dimension of the model;

[1]The label Neumann/free denotes the non-pinned benchmark in which the boundary-active exponential short-end loading is retained. It should not be read as asserting that the operative loading benchmark is the classical homogeneous condition $H'(0) = 0$. The operative content is non-pinning: $H_N(0) = A_N$, so the short end remains stochastic.

2. the deterministic maturity geometry of the curve;
3. the boundary law that selects admissible modes.

The first object concerns the source of randomness. The second concerns the shape of the maturity loading through which that randomness is expressed across the curve. The third concerns the economic regime imposed at the short end.

This distinction explains how a one-factor stochastic model can still generate non-trivial term-structure shapes. A model may be one-factor in randomness while having a richer deterministic maturity geometry. In the IRFR framework, the stochastic dimension and the maturity-mode structure are not the same object. The covariance may be rank-one, while the admissible curve shapes may still include multiple deterministic modes, repeated-root structures, and boundary-induced humps.

The consequence is a different interpretation of model flexibility. Flexibility is not obtained by freely specifying more volatility functions. It arises from the admissible interaction between the rolling-field interior and the boundary condition at the short end.

1.3 The IRFR primitive

The modelling primitive is the instantaneous rolling forward rate

$$x(t,\tau) = F(t, t+\tau), \tau \geq 0$$

Here t denotes calendar time and τ denotes remaining maturity. For a fixed calendar maturity T, remaining maturity satisfies

$$\tau = T - t, d\tau/dt = -1$$

Thus, as calendar time advances, every fixed-maturity contract rolls toward the short end. The field evolves in two ways at once: it is shocked stochastically in calendar time, and it is transported deterministically along the maturity axis.

This transport is not an optional modelling feature. It is a geometric fact of the rolling-maturity representation. Any admissible drift must account for it. A pure mean-reversion term is not enough, because the equilibrium curve itself is being traversed as maturity rolls down. A transport correction is required to preserve consistency between stochastic evolution and the movement of the maturity coordinate.

This is the origin of the IRFR affine structure. The drift is not chosen solely for tractability. It is forced by the requirement that the rolling field remain compatible with equilibrium moment transport. The diffusion is then tied to the same maturity loading. Volatility is therefore not introduced as an independent curve-fitting device; it inherits its shape from the admissible field structure.

The short end, $\tau = 0$, has a special role. In rolling-maturity coordinates it is a true boundary. Every point of the curve is transported toward it. It is also the point at which monetary policy most directly enters the interest-rate system. This makes the boundary mathematically natural and economically meaningful.

The IRFR primitive therefore turns the term-structure problem into a boundary-value problem. The interior law governs admissible rolling-field evolution. The boundary law closes the system.

1.4 HJM: representation versus explanation

The relationship with the Heath–Jarrow–Morton (HJM) framework must be stated carefully [6, 12, 5].

HJM is a broad and powerful arbitrage-free representation of forward-rate dynamics. Given a specification of the forward-rate volatility field, HJM identifies the drift restriction required to preclude arbitrage. In this sense, HJM describes what arbitrage-free forward-rate dynamics may be.

The IRFR framework begins from a different question. It asks whether a particular admissible field structure can be generated from a small set of structural requirements. The point is not to replace HJM as an ambient no-arbitrage representation. An admissible IRFR model should be capable, after translation into standard forward-rate coordinates, of being represented within a suitable HJM-type formulation.

The distinction is therefore not one of arbitrage-free versus non-arbitrage-free modelling. It is one of representation versus explanation.

HJM describes what arbitrage-free forward-rate dynamics may be. IRFR seeks to explain why a particular admissible field structure should arise.

In a standard HJM formulation, the volatility structure is typically specified as an input. The arbitrage-free drift then follows. In the IRFR formulation, the maturity structure of volatility is constrained by the rolling-field geometry. The field's drift, diffusion, maturity loading, and boundary behaviour are not independent modelling choices. They are linked through the admissibility conditions of the IRFR system.

This is why the HJM embedding is not necessarily structure-preserving. The same dynamics may be representable in HJM coordinates, but the reason for their form may become obscure. In the IRFR variables, the short end is a boundary, the maturity coordinate is the field's maturity-domain coordinate, and the admissible modes arise from the rolling-maturity structure. Translated into a conventional HJM representation, these features may appear as a complicated or non-Markovian volatility specification.

Thus HJM should be understood as the ambient arbitrage-free representation, while IRFR is proposed as a boundary-adapted structural explanation inside that broader universe.

This distinction will become particularly important in the repeated-root and boundary-law chapters. A humped volatility loading that appears exogenous in HJM may arise as an admissible boundary mode in IRFR coordinates. The object has not changed. The geometry has.

This also clarifies the role of the Ritchken–Chuang comparison. The claim is not merely that IRFR can generate a humped shape resembling a successful parametric input. The stronger structural point is retrodictive: if humped volatility loadings of the Ritchken–Chuang type have empirical usefulness, then IRFR offers an explanation of why such a maturity profile is natural. The humped form is not imposed as an exogenous volatility ansatz; it is recovered as a boundary-selected admissible mode of the rolling-maturity field. In this sense, the empirical success of humped-volatility specifications becomes indirect evidence for the boundary-field mechanism, while still falling short of a formal equivalence with any particular option-pricing model.

1.5 Boundary laws and monetary policy

The short end of the yield curve is economically special. Central banks operate most directly through short-maturity rates: overnight rates, reserve remuneration, standing facilities, policy corridors, and short-end forward guidance. They do not directly set the entire forward curve. Their influence propagates from the short end into the rest of the maturity domain.

In the IRFR formulation, this observation has a direct mathematical expression. Monetary policy is represented as a boundary law at $\tau = 0$.

A boundary law is not merely an endpoint condition. It is the rule that determines how the field is closed at the short end. Different closures correspond to different economic regimes.

A Neumann/free boundary represents the non-pinned benchmark: the short end remains stochastic and the ordinary exponential short-end loading is retained. A Dirichlet boundary represents an idealised fixed-policy regime, in which short-end stochastic loading is suppressed by the boundary. A Robin boundary represents partial anchoring, in which the short end is influenced by policy but retains residual market stochasticity.

These cases should not be read as literal mechanical descriptions of central-bank operations. They are structural idealisations. Their purpose is to identify how policy regimes select admissible maturity modes.

The Dirichlet case is especially important as a limiting benchmark. Economically, it corresponds to an exaggerated monetary-policy intervention: the short end is fixed. Mathematically, it imposes a local boundary condition on the IRFR field. In the repeated-root branch, this boundary condition generates a linear-exponential humped loading of the form associated with Ritchken–Chuang-type volatility specifications.

This gives the Dirichlet repeated-root case a special interpretive role. In standard HJM finite-dimensional coordinates, a stationary humped volatility specification of this type is usually regarded as non-Markovian. In the IRFR formulation, the same humped structure is generated by a Markovian boundary-value field. What appears as non-Markovian memory in one representation may become a local boundary law in the natural rolling-maturity geometry.

The claim is precise. The model is not asserted to be one-dimensional Markov in the ordinary short-rate sense. It is Markovian only in the field-state sense: an infinite-dimensional boundary-value evolution of the IRFR field, not a finite-dimensional short-rate or low-dimensional HJM state. The recovery of Markovian structure occurs because the state geometry is changed.

The Robin case provides the more realistic policy analogue. It mixes retained short-end stochastic loading with boundary-induced anchoring. The short end is partially anchored rather than fixed. The resulting loading contains both a residual boundary-active component and a boundary-induced component. This gives a structural interpretation to humped or shifted volatility profiles: they reflect the mixture of admissible modes selected by the policy boundary.

In this sense, monetary policy is not modelled as a global curve-wide force. It is modelled as a boundary law whose effects are transmitted through the admissible geometry of the field.

1.6 Comparison with existing literature

A useful comparison is with quadratic term-structure models [1, 8]. Quadratic models enrich the standard affine tradition by allowing bond prices, yields, or pricing equations to depend quadratically on latent state variables. They are powerful because they can generate nonlinear yield dynamics and richer distributional behaviour.

The IRFR framework takes a different route. It does not begin by adding polynomial state dependence. It first changes the geometric primitive from calendar maturity to rolling maturity and asks what drift and loading structures are forced by the resulting boundary-field problem.

In this sense, IRFR is not a substitute for quadratic term-structure models, but a minimalist structural alternative. It shows how non-monotone maturity loadings can arise from admissibility and boundary selection before additional polynomial state machinery is introduced.

1.7 Claims and limitations

The framework is intentionally ambitious, but its claims must be separated.

First, the monograph establishes admissibility. It constructs a class of IRFR dynamics satisfying the proposed structural axioms and boundary closures. This is the mathematical foundation of the work.

Second, it argues for minimality. The claim is not that the framework is the most general arbitrage-free model of interest rates. It is not. Nor is the claim that no other model can reproduce the same empirical shapes. Rather, the claim is that the IRFR construction derives several objects that are often introduced separately: admissible maturity loadings, volatility localisation, boundary anchoring, and policy-induced mode selection. To the extent that these arise from one mechanism, the model is structurally economical.

Third, it motivates a canonicality conjecture. The repeated appearance of the same modes, boundary structures, and humped loadings suggests that the IRFR geometry may be the natural coordinate system for a class of interest-rate phenomena. This is a conjecture, not a theorem. The present work does not prove uniqueness of the IRFR construction among all possible term-structure theories.

Fourth, it proposes an empirical programme. The framework becomes meaningful only if its structural parameters can be related to observed yield curves, volatility profiles, and monetary-policy regimes. The mapping to humped-volatility models provides one bridge, but empirical support must come from calibration and testing rather than assertion.

The pricing chapter introduces one additional diagnostic object for this empirical programme: the Market Signal-to-Noise Ratio (MSNR). MSNR measures the drift signal per unit diffusion noise in the maturity field. Under a risk-neutral measure, the market price of risk is the correction between the physical MSNR and the no-arbitrage MSNR implied by the HJM drift restriction. This gives the pricing adjustment a native IRFR interpretation rather than treating it as an external black-box risk premium.

Fifth, it distinguishes field-level admissibility from pricing no-arbitrage. The affine-closure calculation imposes local coherence on the maturity-domain field: relative stochastic motion across maturities must be compatible with a common loading structure. This condition is no-arbitrage-like in motivation, because it rules out arbitrary inconsistent relative motion along the curve, but it is not itself a self-financing no-arbitrage pricing theorem. Chapter 12 therefore adds a separate bond-pricing layer: the IRFR field is identified with the rolling instantaneous forward-rate curve, the money-market numeraire is introduced, and the HJM martingale restriction is imposed on IRFR-admissible volatility modes. That construction gives IRFR a risk-neutral bond-pricing interpretation, while leaving full derivative-pricing implementation and empirical risk-premium calibration as separate tasks.

The epistemic status of the work can therefore be summarised as follows:

admissibility is proved;

minimality is argued;

canonicality is conjectured;

empirical relevance is to be tested.

This discipline is essential. The purpose of the monograph is not to declare a final theory of interest rates. It is to develop a coherent field formulation and then ask whether the resulting structure is powerful enough to deserve that interpretation.

1.8 Roadmap

The monograph proceeds in five parts.

Part I introduces interest rates as a maturity-domain field. It defines the IRFR primitive, explains rolling-maturity transport, and shows why the short end becomes a boundary. The goal is to establish the geometric language of the theory before the technical machinery is introduced.

Part II develops the axiomatic foundations and transport mechanics. It states the structural axioms, derives the affine drift through KFE-consistent moment transport, and introduces the common maturity loading. It also explains the one-factor covariance structure and the role of signal and noise within a regime.

Part III studies the r-branch and the admissible mode structure. The ratio between the temporal drift-diffusion scale and the maturity-decay scale generates distinct-root and repeated-root regimes. At the critical value r=1, the roots coalesce and the linear-exponential repeated-root mode appears. This mode is central to the structural interpretation of humped volatility.

Part IV develops the boundary laws. Axiom 4 is reinterpreted as the location of economic regime selection. Neumann/free, Dirichlet, and Robin boundary regimes are studied as the non-pinned benchmark, the pinned-policy limit, and the partially anchored regime. The Dirichlet repeated-root case gives the cleanest example of a boundary-selected volatility hump. The Robin case provides the more realistic model of partial policy anchoring.

Part V builds the empirical and modelling bridge. It relates the IRFR boundary modes to Ritchken–Chuang-type humped volatility specifications, compares the minimalist IRFR route with richer nonlinear frameworks such as quadratic term-structure models, translates the results into conventional observable rates, and discusses calibration, testing, and empirical interpretation. It also identifies open problems, including multi-factor IRFR fields, stochastic boundary laws, and no-arbitrage pricing in the presence of a boundary-field short end.

The final principle of the monograph is therefore simple:

Interest rates are maturity-domain fields transported through calendar time and constrained by boundary laws.

The remaining chapters develop this principle mathematically.

Chapter 2

The IRFR Primitive and the Field Representation

2.1 Introduction

The previous chapter introduced the programme of the monograph: to formulate interest-rate dynamics as a maturity-domain field theory in which admissible dynamics arise from rolling-maturity invariance and observable shapes are selected by boundary laws. This chapter develops the primitive object on which that programme is built.

The central modelling choice is to work with the instantaneous rolling forward rate, or IRFR. Instead of describing the forward curve by fixing a calendar maturity T, the IRFR formulation describes the curve by remaining maturity,

$$\tau = T - t$$

The resulting object is

$$x(t, \tau) = F(t, t + \tau), \tau \geq 0,$$

where t is calendar time and τ is time to maturity. The forward curve is therefore represented as a stochastic field over the half-line $[0, \infty)$.

This chapter explains why this change of primitive matters. The move from (t, T) to (t, τ) does not merely relabel the standard forward-rate curve. It exposes the transport structure of the term structure. As calendar time advances, every fixed maturity rolls toward the short end. The point $\tau = 0$ becomes a true boundary of the maturity domain. The model is therefore not just a stochastic process indexed by maturity, but a stochastic field transported through calendar time and closed by a boundary law.

The aim of this chapter is to make that geometry explicit before the axioms and derivations are introduced. In particular, the chapter shows how the coordinate transformation turns maturity into a transport coordinate, separates deterministic roll-down from stochastic evolution, and makes the short end the natural boundary of the field.

2.2 From calendar maturity to rolling maturity

Let $F(t, T)$ denote the instantaneous forward rate observed at calendar time t for maturity date T. The conventional notation treats T as the maturity label. For a fixed T, the process $F(t, T)$ describes the evolution of the forward rate associated with that maturity date.

The IRFR formulation uses remaining maturity instead:

$$\tau = T - t$$

Equivalently,

$$T = t + \tau$$

The instantaneous rolling forward rate is then defined by

$$x(t, \tau) = F(t, t + \tau)$$

For each fixed calendar time t, the map

$$\tau \mapsto x(t, \tau)$$

is the forward curve expressed as a function of remaining maturity. Thus $x(t, \cdot)$ is a curve over the maturity domain, while $t \mapsto x(t, \cdot)$ describes the evolution of that curve through calendar time.

The distinction becomes important when a calendar maturity is held fixed. If T is fixed, then

$$\tau(t) = T - t,$$

and therefore

$$d\tau/dt = -1$$

A forward contract with fixed maturity date does not remain at a fixed point in the maturity domain. It moves deterministically toward the short end as time passes.

For example, consider a contract that has five years remaining maturity at time $t = 0$. At that date it sits at $\tau = 5$. One year later, if the calendar maturity date is unchanged, the same contract has $tau = 4$. It has moved left along the maturity axis even if no new market shock has occurred.

This simple identity is the source of the rolling-field geometry. Calendar time does two things simultaneously: it changes the stochastic state of the curve and it transports each fixed-maturity point along the maturity axis.

2.3 The transport identity

The rolling-maturity representation makes the transport term explicit. Since

$$F(t,T) = x(t, T - t),$$

differentiating along a fixed calendar maturity T gives

$$\frac{d}{dt}F(t,T) = \partial_t x(t,\tau) - \partial_\tau x(t,\tau), \tau = T - t$$

Equivalently, the change in the conventional forward rate along a fixed maturity has two components: the calendar-time change of the rolling field and the deterministic roll-down along the maturity coordinate.

calendar-time evolution: $\partial_t x(t,\tau)$

maturity transport: $-\partial_\tau x(t,\tau)$

A simple example makes the point concrete. Suppose, momentarily, that the curve is linear in remaining maturity:

$$x(t,\tau) = a + b\tau$$

Then $\partial_\tau x(t,\tau) = b$, so the deterministic roll-down contribution is

$$-\partial_\tau x(t,\tau) = -b$$

Thus, even without an explicit calendar-time movement in the field, a fixed calendar maturity experiences deterministic change as it rolls along the sloped curve. What appears as drift in fixed-maturity coordinates may partly be mechanical movement along the curve.

This identity is the first formal indication that the IRFR representation is not merely a relabelling. It separates stochastic evolution in calendar time from deterministic transport in remaining maturity. The transport correction that appears later in the affine drift is a consequence of this coordinate geometry.

The identity also explains why the short end becomes structurally important. Since fixed maturities move in the direction of decreasing τ, all rolling-maturity trajectories are transported toward the boundary $\tau = 0$. The boundary is therefore not appended to the model after the fact; it is already present in the coordinate transformation.

2.4 The forward curve as a maturity-domain field

In the IRFR formulation, the forward curve is treated as a field

$$x : [0,\infty) \times [0,\infty) \to \mathbb{R},$$

where t indexes calendar time and τ indexes remaining maturity. For each t, $x(t,\tau)$ is a function of τ. For each τ, $x(t,\tau)$ evolves stochastically through calendar time.

The term "field" is used in this precise sense. The model is not a collection of unrelated scalar processes. It is a maturity-indexed object whose values across τ are linked by the rolling structure of the curve. The maturity coordinate is not simply a label attached to different instruments. It is the maturity-domain coordinate along which the curve is transported.

The field viewpoint separates two forms of variation.

First, there is stochastic variation in calendar time. This is the usual source of risk in an interest-rate model: the curve changes because the state of the market changes.

Second, there is deterministic variation along the maturity domain. Even in the absence of a new market shock, a fixed calendar maturity rolls toward shorter remaining maturity as time passes. This movement must be reflected in the drift of the field.

A model that ignores this second component risks mixing two distinct effects: genuine stochastic evolution and deterministic maturity transport. The IRFR formulation keeps them separate.

This separation is one reason the framework is able to distinguish the stochastic dimension of the model from the deterministic maturity geometry of the curve. A model may be one-factor in randomness, driven by a single Brownian source, while still possessing a non-trivial maturity-domain structure through its admissible loading functions.

2.5 Local rate exposure and the economic meaning of the IRFR coordinate

The IRFR primitive is not introduced solely because it gives a cleaner mathematical representation. It also has a direct economic interpretation. The local field increment identifies the rate shock at a particular remaining maturity, and hence the local component of risk to which a forward-rate exposure at that maturity is sensitive.

This qualification is important. A forward rate is not itself a full value process. The actual marked-to-market P&L of a traded instrument will generally depend on discounting, annuity factors, contract specification, numeraire choice, and the loading of the instrument on the forward curve. The IRFR increment should therefore be read as the local driver of such an exposure, not as a complete valuation formula.

In conventional notation, the forward rate F(t,T) is attached to a fixed calendar maturity. In rolling-maturity notation, the same rate is represented as $x(t,\tau)$, where τ is the remaining maturity. This makes it possible to locate an infinitesimal forward-rate exposure directly at a point of the maturity domain.

The local field increment

$$dx(t,\tau)$$

therefore represents the instantaneous local rate shock at remaining maturity τ. This is not merely a visual or notational symmetry. It is the connection between the stochastic field and the economically meaningful term-structure exposure. The maturity coordinate describes where the exposure sits on the curve; the calendar-time increment identifies the local movement to which that exposure is sensitive.

This observation is one reason the rolling-maturity representation is structurally useful. The same object that describes the local economics of a forward-rate exposure also supplies the coordinate system in which the curve is transported toward the short-end boundary. The IRFR primitive therefore aligns three elements:

local forward-rate exposure;

rolling maturity transport;

boundary-value geometry.

This alignment is what is meant by the symmetry of the IRFR representation. The symmetry is not merely aesthetic. It means that the economically local object and the geometrically local object coincide. The field coordinate is both the location of a forward-rate exposure and the maturity-domain coordinate along which the curve evolves.

This is precisely what is harder to see in a standard Heath–Jarrow–Morton (HJM) representation. HJM provides the broad arbitrage-free dynamics of the forward curve, but the local economic exposure, the rolling transport geometry, and the short-end boundary condition are not organised around the same coordinate system. In the IRFR formulation they are. This is why the boundary-condition analysis becomes cleaner: the boundary is not introduced as an additional restriction on an already-specified volatility field, but appears as the natural endpoint of the economically meaningful rolling field.

2.6 Transport toward the boundary

The maturity domain of the IRFR field is the half-line

$$\tau \in [0, \infty) .$$

The point $\tau = 0$ corresponds to the instantaneous short end. Every fixed-maturity contract moves toward this point as calendar time advances. Thus the short end is not merely a limiting case of the forward curve. It is the boundary toward which the rolling field is transported.

This observation gives the term-structure problem a boundary-value character. The interior of the maturity domain describes the behaviour of rates with positive remaining maturity. The short end closes the system.

The direction of transport is important. Since

$$d\tau/dt = -1,$$

calendar time carries the field from larger τ toward smaller τ. The boundary at $\tau = 0$ is therefore not an arbitrary endpoint. It is the location at which rolling maturities arrive.

This has two consequences.

First, the drift of the IRFR field must include a transport correction. If the equilibrium curve depends on τ, then a point rolling down the curve experiences a deterministic change even if its stochastic displacement from equilibrium is unchanged. A drift specification that does not account for this movement is not consistent with the rolling-maturity geometry.

Second, the boundary at $\tau = 0$ requires a closure condition. The field cannot be fully specified by its interior law alone. The behaviour at the short end determines which admissible modes are realised.

This is the mathematical origin of the boundary law.

2.7 The economic meaning of the short-end boundary

The boundary at $\tau = 0$ is not only mathematically natural. It is also economically distinguished.

Central banks operate most directly at the short end of the term structure. Policy rates, overnight rates, reserve remuneration, standing facilities, and policy corridors all act first on the shortest maturities. Longer maturities respond through expectations, risk premia, market clearing, and the transmission of policy through the curve. But policy does not usually act by directly fixing every maturity.

For the present chapter, the important point is not yet the classification of possible policy regimes. That will be developed later. The important point is simply that $\tau = 0$ is the mathematically natural location at which short-rate policy enters the maturity-domain field.

Thus the IRFR primitive gives monetary policy a natural location without requiring policy to be modelled as a curve-wide perturbation. Policy enters as a boundary law at $\tau = 0$, and its effects propagate through the admissible maturity geometry of the field.

2.8 Relation to the conventional forward curve

The IRFR field is not a different market object from the conventional instantaneous forward curve. It is a different representation of the same underlying term structure.

Given F(t,T), define

$$x(t, \tau) = F(t, t + \tau)$$

Conversely, given $x(t, \tau)$, the conventional forward rate can be recovered by

$$F(t, T) = x(t, T - t)$$

Thus the two representations are formally equivalent at the level of pointwise curve values. The distinction lies in what the representation makes structural.

The (t,T) representation is natural for no-arbitrage pricing across calendar maturities. It treats each maturity date T as a label. The (t,τ) representation is natural for studying rolling maturity, field transport, and boundary closure. It treats remaining maturity as the domain coordinate.

This is why the IRFR formulation should not be understood as a rejection of standard forward-rate modelling. It is a change of coordinates designed to expose a structure that is otherwise difficult to see.

In particular, the boundary at $\tau = 0$ is explicit in IRFR coordinates. The rolling transport identity is explicit. The maturity-domain field is explicit. These features are present implicitly in any forward-rate model, but they are not always the organising principles of the theory.

2.9 Why the primitive matters

The choice of primitive determines which features of a model are treated as inputs and which are treated as consequences.

If the volatility function is specified directly as a function of maturity, then humped volatility can be introduced whenever it is empirically useful. This may be effective for calibration, but it leaves open the question of why that functional form should arise.

The IRFR programme asks a different question. Given the rolling-maturity field and its boundary closure, which maturity loadings are admissible? Which shapes are selected by boundary conditions? Which volatility profiles arise from the same structure that determines drift and transport?

In this sense, the primitive is not neutral. The IRFR primitive is chosen because it makes the term-structure problem into a field problem with a boundary. Once the field geometry is fixed, the model has less freedom to introduce curve shapes independently.

This is the modelling cost discussed in Chapter 1. The framework gains structural discipline by requiring that curve shape, volatility localisation, maturity loading, and policy anchoring be traced back to identifiable locations in the theory. Some restrictions enter through the interior law. Others enter through the boundary law. The point is to make these restrictions visible.

The IRFR primitive is therefore the first modelling commitment of the monograph.

2.10 Summary

The rolling-maturity representation has four immediate consequences.

First, remaining maturity is a transport coordinate: dτ/dt = -1.

Second, the forward curve is a field over the maturity domain: for each t, $\tau \mapsto x(t, \tau)$ is the curve.

Third, the representation separates stochastic evolution from deterministic roll-down. A local movement in fixed-maturity coordinates contains both calendar-time change and transport along the maturity axis.

Fourth, the short end $\tau = 0$ is a boundary. It is mathematically required for closure and economically distinguished as the natural location of short-rate policy.

These four consequences explain why the IRFR primitive is not merely a convenient notation. It is the coordinate system in which local exposure, roll-down geometry, and boundary analysis are aligned.

2.11 Transition to the axioms

The field representation now motivates the axiomatic structure that follows.

The coordinate transformation has done three things. It has turned remaining maturity into a transport coordinate. It has turned the short end into a boundary. And it has separated stochastic evolution in calendar time from deterministic roll-down along the maturity axis.

The axioms of the next chapter formalise these consequences.

A stochastic model for $x(t, \tau)$ must specify how the field evolves in calendar time. This is the role of the stochastic evolution axiom.

The model must respect the fact that the maturity coordinate is rolling. This is the role of rolling-maturity invariance.

The model must preserve consistency between stochastic evolution and the equilibrium curve as maturities are transported toward the boundary. This is the role of equilibrium moment transport.

Finally, the model must be closed at $\tau = 0$. This is the role of the boundary closure axiom.

Thus the four axioms of the IRFR framework correspond directly to the geometry introduced in this chapter:

stochastic evolution;

rolling-maturity invariance;

equilibrium moment transport;

boundary closure.

The next chapter states these axioms explicitly and explains how they divide the theory into two parts: an interior law determining admissible dynamics, and a boundary law selecting admissible modes.

Chapter 3

The IRFR Axioms

3.1 Introduction

Chapter 2 established the geometric structure of the IRFR field. It showed that the change from calendar maturity T to remaining maturity τ=T-t turns the forward curve into a maturity-domain field, exposes deterministic rolling transport, and identifies the short end $\tau = 0$ as a boundary.

This chapter formalises that geometry through four axioms. Each axiom corresponds to one feature of the IRFR representation:

Geometry established in Chapter 2	**Axiom introduced in Chapter 3**
Stochastic field	Axiom 1: stochastic evolution
Rolling maturity	Axiom 2: rolling-maturity invariance
Transport along equilibrium	Axiom 3: equilibrium moment transport
Boundary at $\tau = 0$	Axiom 4: boundary closure

The axioms are not arbitrary modelling assumptions. They are the formal counterparts of the geometry already established. Axiom 1 gives the field local stochastic evolution. Axiom 2 prevents the model from assigning unrelated dynamics to two contracts with the same remaining maturity observed at different calendar dates. Axiom 3 requires the expected dynamics to be consistent with roll-down along the equilibrium curve. Axiom 4 closes the field at the short-end boundary.

Without Axiom 2, the coordinate τ would lose its meaning as a maturity-domain coordinate. It would become only a relabelling of calendar maturities. The purpose of the axioms is therefore to preserve the interpretation of $x(t, \tau)$ as a rolling field rather than a collection of unrelated forward-rate processes.

The first three axioms determine the interior dynamics of the field. They restrict how the IRFR may evolve away from the boundary. The fourth axiom determines how the system is closed at $\tau = 0$, and therefore where economic regime enters the model.

Axioms 1–3 determine admissible interior dynamics;

Axiom 4 determines the boundary law.

This division gives mathematical form to the programme stated in Chapter 1: dynamics are determined by invariance, while observable shapes are selected by boundary laws.

3.2 Why axioms?

The IRFR framework uses axioms to make the modelling cost visible. In term-structure models, assumptions often enter through several different choices: the state variable, the factor structure, the drift, the volatility function, the long-run mean, the policy rule, and the calibration procedure. These choices may be individually reasonable, but their structural role is not always explicit.

The axioms assign each restriction to a specific location in the theory. The stochastic character of the curve enters through Axiom 1. The rolling geometry of remaining maturity enters through Axiom 2. The compatibility between stochastic evolution and deterministic roll-down enters through Axiom 3. The economic regime at the short end enters through Axiom 4.

This is the sense in which the framework seeks internal consistency rather than mere parsimony. The aim is not simply to use fewer assumptions. It is to ensure that curve shape, volatility structure, maturity loading, and boundary behaviour can be traced back to governing principles.

The axioms should therefore be read as structural commitments. They define the class of admissible IRFR fields considered in this monograph. Later chapters derive their consequences.

3.3 Axiom 1: stochastic evolution

The first axiom states that the IRFR field evolves stochastically through calendar time.

For each remaining maturity $\tau > 0$, the local evolution of the field is assumed to admit a diffusion representation,

$$dx(t, \tau) = -f(x, t, \tau)dt + \sqrt{2}g(x, t, \tau)dW_t,$$

where W_t is a Brownian motion, f is the IRFR drift-operator coefficient, g is the IRFR diffusion-scale coefficient, and the actual SDE drift is -f.

This is the local stochastic evolution axiom.

The axiom does not yet specify the functional form of f or g. It only states that, locally in calendar time, the IRFR field evolves as a diffusion. The later axioms restrict the admissible relationship between the drift operator, the diffusion scale, and the maturity coordinate.

The use of a single Brownian driver reflects the one-factor version of the theory developed in the first part of the monograph. This does not mean that the maturity structure is one-dimensional. It means that the source of randomness is one-dimensional. The way this randomness is expressed across maturity is governed by the admissible maturity loading.

In later chapters, the diffusion scale will be written using a maturity-loading convention, for example through a function H. The precise variance-rate convention is fixed when the Kolmogorov forward equation is introduced. For now, the important point is simply that the field has local stochastic dynamics.

3.4 Interpretation of Axiom 1

Axiom 1 gives the IRFR field its local risk structure.

At each remaining maturity τ, the local field increment

$$dx(t, \tau)$$

represents the local rate shock to which an infinitesimal forward-rate exposure at remaining maturity τ is sensitive. The full mark-to-market P&L of a traded instrument will generally depend on discounting, annuity effects, contract-specific loading, and numeraire choice. The IRFR increment is therefore not itself a complete valuation object. It identifies the local field component driving that exposure.

This is the economic role of Axiom 1. It connects the stochastic field to local forward-rate risk without requiring the forward rate itself to be treated as a full value process.

The actual SDE drift, -f, describes the expected local movement of the field. The diffusion-scale term g describes local uncertainty, with $\sqrt{2g}$ as the conventional volatility coefficient. The Brownian driver represents the common stochastic shock. The maturity dependence of the response to that shock is not yet specified; it will be constrained by the later admissibility conditions.

This is deliberate. The IRFR framework does not begin by choosing a volatility curve. It begins by specifying that the field evolves stochastically, then uses rolling invariance, moment transport, and boundary closure to determine which structures are admissible.

3.5 Axiom 2: rolling-maturity invariance

The second axiom states that the structural form of the IRFR dynamics is invariant under the rolling-maturity representation.

The remaining maturity coordinate is

$$\tau = T - t$$

For a fixed calendar maturity T,

$$d\tau/dt = -1$$

Thus, as calendar time advances, a fixed-maturity point moves along the maturity axis. A model of the IRFR field must be consistent with this movement.

Axiom 2 imposes that, when expressed in rolling-maturity coordinates, the structural form of the local dynamics depends on remaining maturity τ, rather than on the calendar maturity date T itself. Equivalently, within a given regime, the same remaining maturity is governed by the same structural law across calendar time.

For example, without Axiom 2, the model could assign different structural dynamics to two contracts that both have τ=5, merely because they are observed at different calendar dates. This would break the interpretation of τ as a maturity-domain coordinate rather than a calendar label.

A one-year remaining maturity should be governed by the one-year structural law, regardless of the calendar date on which it is observed, except insofar as the state or regime of the field has changed. This is the invariance principle underlying the IRFR construction.

3.6 Interpretation of Axiom 2

Rolling-maturity invariance is the geometric heart of the model.

Without it, the forward curve could still be written as $x(t, \tau)$, but the notation would impose no discipline. The model could assign arbitrary dynamics to each point of the maturity domain. Axiom 2 prevents this by giving the maturity coordinate intrinsic meaning.

The point τ=5 represents five years of remaining maturity, not a particular calendar maturity date. If a contract has five years remaining today, and another contract has five years remaining at a later calendar date, the same structural law applies to both, within the same regime.

This is what allows the model to speak meaningfully about maturity loading. If the response of the curve to a stochastic shock depends on τ, then that dependence is not merely a label attached to a particular maturity date. It is part of the field's deterministic maturity geometry.

Axiom 2 is therefore the bridge between the economic object and the field representation. It is also the first step toward boundary analysis. Once τ is the maturity-domain coordinate, $\tau = 0$ is the boundary of that domain.

3.7 Axiom 3: equilibrium moment transport

The third axiom imposes consistency between stochastic evolution and deterministic roll-down.

Let

$$m(\tau)$$

denote an equilibrium mean curve for the IRFR field. Since m is a function of remaining maturity, a fixed calendar maturity rolling through time also moves along this equilibrium curve.

For a fixed T, remaining maturity is

$$\tau = T - t$$

The equilibrium value associated with that fixed maturity is therefore

$$m(T - t)$$

Differentiating with respect to calendar time gives

$$\frac{d}{dt} m(T - t) = -m'(\tau)$$

The sign is important. As calendar time advances, the maturity coordinate decreases. A fixed-maturity point rolls toward the short end.

Axiom 3 requires the first moment of the stochastic dynamics to be compatible with this deterministic roll-down. Informally,

expected stochastic evolution must be consistent with transport along $m(\tau)$.

This condition is what later forces the drift structure to contain two components: a displacement term measuring deviation from the equilibrium curve and a transport correction reflecting movement along the maturity axis.

The detailed derivation is deferred to the next chapter. For now, the point is structural. If the field has an equilibrium curve, and maturities roll toward the boundary, then the expected motion of the field must respect the fact that the equilibrium curve is itself being traversed.

3.8 Running example: a sloped equilibrium curve

A simple running example makes the role of transport visible.

Suppose the equilibrium curve is locally linear,

$$m(\tau) = a + b\tau$$

For a fixed calendar maturity T, the equilibrium value along the rolling maturity is

$$m(T-t) = a + b(T-t)$$

Differentiating gives

$$\frac{d}{dt}m(T-t) = -b$$

Thus, even if the field is exactly on the equilibrium curve and no new stochastic shock occurs, the value associated with a fixed maturity changes mechanically as the contract rolls along the curve.

If b>0, the equilibrium curve is upward sloping in τ, and the fixed-maturity point rolls downward as time passes. If b<0, the roll-down effect has the opposite sign.

The flat case b=0 is useful as a limiting benchmark: there is then no roll-down correction. But precisely because the flat case hides the transport effect, the sloped case reveals why the correction is needed.

A model that omits this deterministic roll-down will misidentify part of the curve's mechanical movement as stochastic drift. Axiom 3 prevents this by requiring the expected dynamics to respect the transport of the equilibrium curve.

This example is deliberately simple, but it captures the essential point. In a rolling-maturity field, the drift must be compatible not only with stochastic mean reversion, but also with motion along the maturity axis.

3.9 Interpretation of Axiom 3

Axiom 3 is where the coordinate geometry begins to constrain the dynamics.

A standard local diffusion can specify drift and volatility at each state. But an IRFR field is not merely a collection of local diffusions. A point associated with a fixed calendar maturity is moving through the maturity domain. Its expected evolution must therefore include both stochastic displacement and deterministic roll-down.

Suppose the field is exactly on the equilibrium curve. As calendar time passes, a fixed-maturity point should remain consistent with the equilibrium curve as its remaining maturity declines. If the drift does not account for this movement, the model creates an inconsistency between the probabilistic evolution of the rate and the deterministic geometry of remaining maturity.

This implies that the affine drift-operator structure derived later must contain a transport correction. The correction is not an arbitrary adjustment. It is required by the rolling-maturity geometry.

Axiom 3 therefore connects the probabilistic dynamics of the field with its deterministic maturity geometry.

3.10 Axiom 4: boundary closure

The fourth axiom states that the IRFR field must be closed at the short-end boundary,

$$\tau = 0$$

The maturity domain is the half-line $[0, \infty)$. The point $\tau = 0$ is the instantaneous short end. Since fixed-maturity points are transported toward this boundary, the interior dynamics alone do not fully determine the field. A boundary condition is required.

Axiom 4 is the boundary closure axiom.

At this stage, the axiom should be understood abstractly. It does not prescribe a single universal boundary condition. It states that an admissible IRFR model must specify how the field behaves at the short end.

The distinction is crucial:

the need for boundary closure is structural;

the form of boundary closure is regime-dependent.

In the original non-policy benchmark, the boundary may be interpreted as a natural or positivity-preserving closure. In the monetary-policy extension, different policy regimes correspond to different boundary laws. These cases will be studied later. The present chapter only identifies the boundary as the location at which closure enters.

Thus Axiom 4 is the point at which economic regime enters the model.

3.11 Interpretation of Axiom 4

Axiom 4 completes the field problem.

Axioms 1–3 determine the admissible interior dynamics. They specify that the field evolves stochastically, respects rolling-maturity invariance, and transports moments consistently along the equilibrium curve. But they do not by themselves determine how the field is closed at the short end.

The boundary condition performs this selection. It is the closure rule for the maturity-domain field, not merely another local modelling assumption. Different closures can select different admissible modes, and therefore different observable curve and volatility shapes.

Economically, this is natural. The short end is where monetary policy acts most directly. Mathematically, it is the boundary of the field. The boundary law therefore provides the link between policy regime and term-structure shape.

The later boundary chapters will study Neumann/free closure, Dirichlet fixed-boundary closure, and Robin partial-anchoring closure. The present chapter does not analyse these cases. It only establishes where they enter the theory.

3.12 Interior law versus boundary law

The four axioms separate the model into two components.

The first component is the interior law, governed by Axioms 1–3. It determines how the field evolves away from the boundary, how the maturity coordinate constrains the dynamics, and how the drift must remain compatible with equilibrium transport.

The second component is the boundary law, governed by Axiom 4. It determines how the field is closed at the short end and which admissible behaviour is selected.

This separation allows the model to distinguish between the set of dynamically admissible structures and the economic regime that selects from them.

Axioms 1–3: admissible interior dynamics

Axiom 4: boundary selection

This is the technical version of the monograph's organising principle: invariance determines the admissible dynamics, while boundary laws select observable shapes.

3.13 Summary and transition

The IRFR framework is built from four structural axioms. Axiom 1 gives the field local stochastic evolution. Axiom 2 imposes rolling-maturity invariance. Axiom 3 requires equilibrium moment transport. Axiom 4 closes the field at the short-end boundary.

Together they organise the model as a boundary-value problem for a stochastic maturity-domain field. The first three axioms determine the interior law: they restrict the admissible drift, diffusion, and maturity loading. The fourth axiom determines the boundary law: it selects which admissible behaviour is realised at the short end.

The next chapter turns the axioms into equations. Axiom 1 provides the local diffusion form. Axiom 2 imposes rolling-maturity invariance. Axiom 3 requires consistency of first moments under deterministic roll-down. Together, these conditions force the affine drift-operator structure and the common maturity loading that underlie the IRFR model.

The Kolmogorov forward equation provides the mechanism for this derivation. It links the local drift and diffusion of the field to the evolution of the probability density and its moments. When combined with rolling-maturity transport, it yields the admissibility conditions that govern the interior law.

The next chapter therefore studies the Kolmogorov forward equation, affine closure, and moment transport. It is there that the axioms stated here become calculable restrictions on the form of the IRFR dynamics.

Chapter 4

The Kolmogorov Forward Equation and Affine Closure

4.1 Introduction

Chapter 3 stated the four axioms of the IRFR framework. The first three axioms determine the interior law of the field: local stochastic evolution, rolling-maturity invariance, and equilibrium moment transport. The fourth axiom closes the field at the short-end boundary.

This chapter turns the interior axioms into equations. Its purpose is to show how the local stochastic specification of the IRFR field becomes a calculable restriction on drift, diffusion, and maturity loading.

The required tool is the Kolmogorov forward equation. The KFE links the local drift and diffusion of a stochastic process to the evolution of its probability density and moments. In the present setting, it is the mechanism through which the axioms become admissibility conditions.

The main result is that the IRFR drift operator cannot be chosen freely. Once the field is required to evolve as a diffusion, respect rolling maturity, and transport its equilibrium moments consistently, the drift structure must combine two components: a displacement term relative to the equilibrium curve and a transport correction generated by deterministic roll-down along the maturity axis.

The IRFR field is

$$x(t, \tau) = F(t, t + \tau), \tau \geq 0$$

For each remaining maturity τ, the local stochastic evolution is written as

$$dx(t, \tau) = -f(x, t, \tau)dt + \sqrt{2}g(x, t, \tau)dW_t,$$

where f is the IRFR drift operator coefficient and g is the IRFR diffusion-scale coefficient. The actual SDE drift is -f. In conventional SDE notation, the volatility coefficient would be $\sigma{=}\sqrt{}2$g. This chapter focuses on the interior law.

Boundary closure is not yet imposed. Boundary conditions return later, once the admissible interior dynamics have been derived.

4.2 From axioms to a density equation

Axiom 1 states that, at each remaining maturity τ, the IRFR field evolves locally as a diffusion,

$$dx(t,\tau) = -f(x,t,\tau)dt + \sqrt{2}g(x,t,\tau)dW_t$$

For a fixed τ, let

$$p(x,t;\tau)$$

denote the probability density of the IRFR value $x(t,\tau)$. The Kolmogorov forward equation describes how probability mass flows over the state variable under drift and diffusion.

With the convention

$$dx = -f(x,t,\tau)dt + \sqrt{2}g(x,t,\tau)dW_t,$$

the density satisfies

$$\frac{\partial p}{\partial t} = \partial/\partial x[f(x,t,\tau)p] + \partial^2/\partial x^2[g^2(x,t,\tau)p]$$

It is useful to write the instantaneous variance rate as

$$v(x,t,\tau) = 2g^2(x,t,\tau)$$

Then the KFE becomes

$$\frac{\partial p}{\partial t} = \partial/\partial x[fp] + 1/2\partial^2/\partial x^2[vp]$$

This is the sign convention used throughout the moment calculations. Since the local SDE drift is $-f$, the first-order KFE term appears as $\partial_x[fp]$, not as $-\partial_x[fp]$. The integrations by parts below use this convention consistently; the transport correction enters through the affine form of f, rather than as a separate sign change in the state-density operator.

This equation is local in the state variable x, but its coefficients may depend on the maturity coordinate τ. The role of Axioms 2 and 3 is to restrict that τ-dependence. The KFE is therefore the bridge between stochastic evolution and moment transport.

4.3 Moment equations

The first moment of the IRFR density is

$$M_1(t,\tau) = E[x(t,\tau)] = \int xp(x,t;\tau)dx$$

Differentiating under the integral sign and using the KFE gives

$$\frac{\partial M_1}{\partial t} = \int x(\frac{\partial p}{\partial t})dx$$

Substituting the KFE and integrating by parts, assuming boundary terms vanish under the admissibility conditions, yields

$$\frac{\partial M_1}{\partial t} = -E[f(x,t,\tau)]$$

The integration by parts step moves derivatives from the density onto the test function x. Since the second derivative of x is zero, the diffusion contribution drops out of the first-moment equation, leaving the actual drift contribution, -E[f], in expectation.

Thus the expected drift of the SDE is -E[f]. This is the first point at which the axioms impose structure: Axiom 3 requires this moment evolution to be consistent with deterministic roll-down along the equilibrium curve.

Let the equilibrium mean curve be

$$m(\tau)$$

If the field is in equilibrium, then

$$M_1(t,\tau) = m(\tau)$$

For a fixed calendar maturity T, the remaining maturity is τ=T-t. The equilibrium value associated with that fixed maturity is therefore

$$m(T-t)$$

As shown in Chapter 3,

$$\frac{d}{dt}m(T-t) = -m'(\tau)$$

Equivalently, the calendar-time evolution of the rolling field must include a transport correction associated with the derivative of m with respect to τ. This is the moment-transport requirement.

4.4 The transport correction

To see why a transport correction is unavoidable, return to the running example

$$m(\tau) = \alpha + \beta\tau$$

For a fixed calendar maturity T,

$$m(T-t) = \alpha + \beta(T-t),$$

and hence

$$\frac{d}{dt}m(T-t) = -\beta$$

Since

$m'(\tau)$=β,

the roll-down contribution is

-$m'(\tau)$=-β.

Thus, even when the field lies exactly on the equilibrium curve, the fixed-maturity value changes mechanically as remaining maturity declines.

If the drift were specified only as a displacement from equilibrium, it would vanish on the equilibrium curve. That would be inconsistent with roll-down whenever $m'(\tau)$ is non-zero. The drift must therefore include a term that accounts for movement along the equilibrium curve.

The precise sign of this term depends on whether one refers to the IRFR operator coefficient f or to the actual SDE drift -f. The fixed-maturity roll-down contribution is -$m'(\tau)$, while the compensating term in the fixed-τ rolling-field drift appears with the opposite sign. Since the actual drift is -f, the corresponding term inside f appears with a minus sign. In the convention used here, the admissible IRFR drift operator takes the form

$$f(x,t,\tau) = a^2[x - m(\tau)] - m'(\tau)$$

The term -$m'(\tau)$ inside f produces the $+m'(\tau)$ transport correction in the actual SDE drift -f. It is not an arbitrary adjustment. It is required by the coordinate geometry.

4.5 Affine closure of the drift operator

The next question is why the drift operator should be affine in x.

The first-moment equation is

$$\frac{\partial M_1}{\partial t} = -E[f(x,t,\tau)]$$

Closure of the first moment requires E[f(x,t,τ)] to depend only on $M_1(t,\tau)$, rather than on higher moments of the density. This holds for all admissible densities when f is at most affine in x. A nonlinear drift operator would generally cause the first-moment equation to depend on higher moments, breaking first-moment closure.

Accordingly, consider the affine form

$$f(x,t,\tau) = A(\tau)x + B(\tau),$$

where $A(\tau)$ and $B(\tau)$ are functions of remaining maturity. Then

$$E[f(x,t,\tau)] = A(\tau)M_1(t,\tau) + B(\tau)$$

At equilibrium, $M_1(t,\tau)$=$m(\tau)$. The moment-transport condition requires the actual expected drift, -E[f], to match the required transport of the equilibrium curve. This imposes a relation between $A(\tau)$, $B(\tau)$, and $m(\tau)$.

Choosing the displacement scale as a^2, the affine drift can be written as

$$f(x,t,\tau) = a^2[x - m(\tau)] - m'(\tau)$$

The sign convention used here is discussed explicitly in Section 4.7.

Equivalently,

$$f(x,t,\tau) = a^2x - a^2m(\tau) - m'(\tau)$$

Thus, in the affine representation,

$A(\tau) = a^2$,

and

$B(\tau) = -a^2m(\tau) - m'(\tau)$.

The affine drift operator is therefore not introduced merely for tractability. It is the closure form compatible with first-moment transport under the IRFR geometry.

4.6 Interpretation of the affine drift

The IRFR drift operator

$$f(x,t,\tau) = a^2[x - m(\tau)] - m'(\tau)$$

contains two distinct terms. When inserted into the local evolution equation, the actual SDE drift is

$$-f(x,t,\tau) = -a^2[x - m(\tau)] + m'(\tau)$$

This displayed form makes explicit that the actual SDE drift contains the standard mean-reverting displacement term together with the fixed-τ transport correction.

The first term,

a^2[x-$m(\tau)$],

measures displacement from the equilibrium curve. Through the actual drift -f, it produces the mean-reverting contribution $-a^2[x - m(\tau)]$.

The second term,

-$m'(\tau)$,

is the transport component inside f. Through the actual drift -f, it produces the $+m'(\tau)$ transport correction required by fixed-τ rolling-field dynamics.

These two terms have different origins. The displacement term is probabilistic: it describes the local tendency of the stochastic field relative to equilibrium. The transport term is geometric: it is forced by the rolling coordinate.

The running example makes the distinction visible. If

$$m(\tau) = \alpha + \beta\tau,$$

then

$m'(\tau)$=β.

The transport component is constant. It vanishes only when $\beta = 0$, that is, when the equilibrium curve is flat. In a sloped equilibrium curve, the correction is unavoidable.

This is the first major payoff of the IRFR formulation. The drift operator is not simply chosen. It is decomposed into a stochastic displacement component and a deterministic transport component, each with a clear structural origin.

4.7 Sign conventions and mean reversion

The affine drift operator written above uses the convention

$$f(x, t, \tau) = a^2[x - m(\tau)] - m'(\tau)$$

The local stochastic evolution is written with -f as the actual drift,

$$dx(t, \tau) = -f(x, t, \tau)dt + \sqrt{2}g(x, t, \tau)dW_t$$

Thus f should not be read directly as the SDE drift coefficient. The actual drift is -f. Under this convention, the displacement component a^2[x-$m(\tau)$] inside f produces the standard mean-reverting contribution $-a^2[x - m(\tau)]$ in the SDE.

We retain the IRFR convention because it is the convention used in the admissibility algebra and in the later boundary-condition analysis. It also allows the structural comparison between f and g to be written directly. The essential point is the consistency of the sign convention across the local evolution equation, the KFE, the moment-transport condition, the equilibrium density, and the boundary laws.

For this reason, the affine form should be read as part of the complete IRFR operator convention. Later chapters retain this convention throughout. Whenever the formula is translated into a standard SDE notation, the corresponding drift coefficient is -f.

4.8 Variance and common maturity loading

While the drift is determined by first-moment transport, the diffusion structure enters through second-moment consistency.

Axiom 1 allows the local diffusion scale g to depend on the maturity coordinate. Axiom 2 prevents this dependence from being arbitrary across calendar maturities. Axiom 3 links the diffusion structure to the same maturity geometry that governs the drift.

The one-factor specification means that all maturities share a common Brownian driver,

dW_t.

Therefore the covariance between two maturity points is determined by the product of their actual diffusion coefficients. If the diffusion-scale coefficient has the form

$$g(\tau) = aH(\tau),$$

so that the conventional volatility coefficient is $\sqrt{2}g(\tau)$=$\sqrt{2}aH(\tau)$, then the instantaneous covariance between maturities τ_1 and τ_2 is

$$\mathrm{Cov}(dx(t,\tau_1), dx(t,\tau_2)) = 2a^2 H(\tau_1)H(\tau_2)dt$$

The normalisation with $\sqrt{2}$ in the SDE is chosen for later algebraic convenience. The economically relevant object is the maturity loading $H(\tau)$, which determines how the common stochastic shock is expressed across the curve.

This covariance kernel is rank one. The stochastic dimension is one, but the deterministic maturity loading $H(\tau)$ can be non-trivial.

This distinction will matter later. It is what allows a one-factor stochastic model to generate structured maturity profiles without introducing additional Brownian factors.

4.9 Moment closure and admissibility

The KFE generates equations for all moments of the density. The first moment gives the drift condition. The second moment introduces the variance-rate structure.

Let

V(t,τ)=Var[$x(t,\tau)$].

Under the diffusion specification, the evolution of the second moment depends on both the actual drift -f and the variance rate

$$v(x,t,\tau) = 2g^2(x,t,\tau)$$

For the IRFR model to remain admissible, the variance structure must close consistently with the same maturity loading that appears in the drift.

This is where coefficient matching enters. Coefficient matching requires that the moment equations generated by the KFE hold for all admissible states. This imposes algebraic restrictions on the drift operator and variance coefficients. In practice, this means equating the τ-dependent coefficients generated by the moment equations with the proposed functional form of the maturity loading.

Conceptually, coefficient matching asks whether the proposed drift operator and diffusion scale can satisfy the KFE moment equations simultaneously. If they can, the model is admissible. If not, the chosen maturity loading is incompatible with the IRFR axioms.

Thus admissibility is not imposed by assertion. It is tested through the KFE. The affine drift operator and the common maturity loading are the first outputs of that test.

4.10 Relation to HJM

From the HJM perspective, the preceding construction can be read as identifying a restricted admissible subclass of forward-rate dynamics.

The distinction is that the diffusion or volatility field has not been freely specified and then paired with the corresponding arbitrage-free drift. Instead, the admissibility conditions restrict the drift operator and maturity loading together. The KFE and moment-transport equations act as a consistency filter.

This is the technical version of the representation-versus-explanation distinction introduced in Chapter 1. An IRFR model may be representable in HJM coordinates, but the reason for its form is clearest in rolling-maturity field coordinates.

The HJM representation describes what arbitrage-free forward-rate dynamics may be. The IRFR derivation explains why a particular boundary-adapted field structure is admissible.

4.11 Summary

This chapter has turned the first three IRFR axioms into the beginning of a calculable interior law.

Axiom 1 gives the local diffusion. Axiom 2 imposes rolling-maturity invariance. Axiom 3 requires equilibrium moment transport. Together, these conditions force the drift operator to take an affine transport-compatible form,

$$f(x, t, \tau) = a^2[x - m(\tau)] - m'(\tau)$$

The first term measures displacement from equilibrium and generates mean reversion through the actual drift -f. The second term is the transport component inside f, generating the required $+m'(\tau)$ correction in the actual drift.

The chapter has also introduced the role of the common maturity loading in the diffusion structure. In the one-factor case, all maturities share the same Brownian shock, but the deterministic loading determines how that shock is expressed across the curve.

The main message is that the interior law is not freely chosen. It is constrained by the interaction between stochastic evolution, rolling maturity, and moment transport.

4.12 Transition to admissible modes

The next step is to determine which maturity loadings are admissible.

The affine drift operator identifies the structure of local mean evolution, but the full IRFR model also requires the diffusion and variance structure to satisfy the KFE consistently. This leads to coefficient matching and to the emergence of admissible maturity modes.

The next chapter develops this admissibility calculation. It shows how the KFE restricts the maturity loading, how the r-branch arises, and why the repeated-root case later plays a central role in the boundary-law interpretation of humped volatility.

Chapter 5

Admissible Maturity Modes and the r-Branch

5.1 Introduction

Chapter 4 derived the affine drift operator of the IRFR field. Under the convention

$$dx(t,\tau) = -f(x,t,\tau)\,dt + \sqrt{2}g(x,t,\tau)\,dW_t,$$

the Kolmogorov forward equation and first-moment transport imply the drift-operator form

$$f(x,t,\tau) = a^2[x - m(\tau)] - m'(\tau)$$

The actual SDE drift is therefore

$$-f(x,t,\tau) = -a^2[x - m(\tau)] + m'(\tau),$$

which contains standard mean reversion toward the equilibrium curve together with the fixed-τ transport correction.

Chapter 4 also introduced the diffusion-scale coefficient g, with instantaneous variance rate

$$v(x,t,\tau) = 2g^2(x,t,\tau)$$

The task of the present chapter is to determine which maturity loadings are admissible. The drift structure alone does not complete the model. The diffusion structure must also satisfy the Kolmogorov forward equation consistently with the same maturity-domain geometry. This requirement leads to coefficient matching and to a family of admissible maturity modes.

The key result is that the admissible loading is not arbitrary. It is governed by a second-order maturity equation whose solutions divide into branches. These branches are indexed by a dimensionless parameter r, which measures the relationship between the temporal drift-diffusion scale and the maturity-decay scale.

The chapter develops three ideas:

coefficient matching restricts the maturity loading;

the resulting maturity equation generates an r-branch;

the repeated-root case $r = 1$ produces the linear-exponential mode.

The repeated-root case will later become central to the boundary-law interpretation of humped volatility.

5.2 From affine closure to admissibility

The affine drift operator derived in Chapter 4 is

$$f(x, t, \tau) = a^2[x - m(\tau)] - m'(\tau)$$

This form is required by first-moment closure and equilibrium transport. It ensures that the expected dynamics of the field are compatible with deterministic roll-down along the equilibrium curve.

But admissibility requires more than first-moment consistency. The full probability density must evolve consistently under the KFE. In particular, the second moment and variance rate must close in a way compatible with the same maturity loading.

The diffusion-scale coefficient g determines the local variance rate,

$$v(x, t, \tau) = 2g^2(x, t, \tau)$$

In the one-factor model, all maturities share the same Brownian shock, so the covariance between maturity points is generated by the product of their maturity loadings. If

$$g(\tau) = aH(\tau),$$

then the covariance kernel is

$$\mathrm{Cov}(dx(t, \tau_1), dx(t, \tau_2)) = 2a^2 H(\tau_1)H(\tau_2)\, dt$$

Thus the stochastic dimension is one, but the deterministic maturity loading $H(\tau)$ may have non-trivial shape.

The question is therefore:

which functions $H(\tau)$ are compatible with the IRFR axioms?

This is the admissibility problem.

5.3 Coefficient matching

Coefficient matching is the algebraic test for admissibility.

The KFE generates moment equations. The affine drift operator determines how the first moment evolves. The diffusion-scale coefficient determines how the second moment and variance rate evolve. For a proposed maturity loading to be admissible, these equations must be mutually consistent for all admissible states in the model class.

In practice, this means that the τ-dependent coefficients generated by the moment equations must match the proposed functional form of the maturity loading. If the terms generated by the KFE cannot be expressed in the same loading basis, the proposed structure is not closed under the dynamics.

This is the point at which the IRFR framework becomes restrictive. One cannot choose an arbitrary maturity loading and then regard it as admissible. The loading must survive the coefficient-matching test.

The result is a maturity-domain equation for the admissible loading $H(\tau)$. In canonical form, this equation may be written as

$$H''(\tau) + \lambda H'(\tau) + \mu H(\tau) = 0,$$

where λ and μ are constants determined by the drift-diffusion scale and the maturity-domain normalisation.

The exact normalisation will be fixed below through the r-parameter. At this stage, the important structural point is that admissible maturity loadings are not arbitrary functions. They are modes of a second-order linear equation on the maturity domain.

5.4 The characteristic equation

The admissibility equation admits exponential trial solutions of the form

$$H(\tau) = e^{-k\tau}$$

Substituting this into

$$H''(\tau) + \lambda H'(\tau) + \mu H(\tau) = 0$$

gives

$$k^2 - \lambda k + \mu = 0$$

Thus the admissible decay rates are the roots of the characteristic equation.

In the generic case, the two roots are distinct:

$$k_1 \neq k_2$$

The admissible loading is then

$$H(\tau) = C_1 e^{-k_1 \tau} + C_2 e^{-k_2 \tau}$$

For admissible decaying modes on the half-line, the relevant roots have positive real part. In the real distinct-root case, this means

$k_1 > 0$, $k_2 > 0$.

The constants C_1 and C_2 are not determined by the interior equation alone. They are selected by normalisation and boundary closure. This is the first point at which the distinction between interior law and boundary law becomes concrete:

the interior law determines the admissible modes;

the boundary law selects their realised combination.

5.5 The r-parameter

It is useful to parameterise the characteristic equation by a dimensionless quantity r. The role of r is to measure the relative strength of the temporal drift-diffusion scale and the maturity-domain decay scale.

Equivalently, r controls the discriminant of the characteristic equation. In the normalisation used in the remainder of the monograph, r is chosen so that the discriminant is proportional to r-1. With the canonical characteristic equation

$$H''(\tau) + \lambda H'(\tau) + a^2 H(\tau) = 0,$$

this normalisation may be written

$$r = \frac{\lambda^2}{a^2}.$$

Here λ denotes the coefficient in the maturity-domain characteristic equation, not the market price of risk introduced later in Chapter 12. Thus

$$r \neq 1$$

corresponds to distinct roots, while

$$r = 1$$

is precisely the repeated-root case.

This is the key mathematical role of r. It is not introduced as an additional modelling factor. It is a dimensionless way of describing which branch of the admissible maturity equation is realised.

The three relevant cases are therefore:

$r > 1$: distinct real roots under the chosen normalisation;

$r = 1$: repeated root.

$r < 1$: a non-generic branch whose admissibility depends on the chosen real/complex root convention and boundary domain.

The detailed interpretation of the r < 1 branch will therefore depend on the exact normalisation and boundary regime. The repeated-root case r = 1, however, has an invariant significance: it is the point at which the two admissible decay modes coalesce.

5.6 Distinct-root branch

When the characteristic roots are distinct, the admissible loading is

$$H(\tau) = C_1 e^{-k_1 \tau} + C_2 e^{-k_2 \tau}, k_1 \neq k_2$$

This branch represents a mixture of two maturity-decay modes. The stochastic dimension of the model remains one, because all maturities are driven by the same Brownian shock. But the deterministic maturity geometry may contain more than one decay scale.

This distinction is essential. A model can be one-factor in randomness while still having a structured maturity-domain response. The covariance kernel remains rank one, but the loading through which the common shock is expressed across the curve can have non-trivial shape.

The distinct-root branch is therefore the generic admissible case. Boundary conditions and normalisation determine which linear combination of the two modes is realised.

5.7 Repeated-root branch

The repeated-root case occurs when

$$k_1 = k_2 = k$$

At this point, the two-exponential representation degenerates. The second independent solution is no longer a distinct exponential. Instead, the admissible loading takes the form

$$H(\tau) = (C_1 + C_2 \tau) e^{-k\tau}$$

For admissible decaying modes on the half-line, we take

$$k > 0$$

The second term,

$$\tau e^{-k\tau},$$

is the important one. It is not a simple exponential decay. It vanishes at the short end, rises to an interior maximum, and then decays to zero as ($\tau \to \infty$).

Indeed, for

$$h(\tau) = \tau e^{-k\tau},$$

we have

$$h'(\tau) = e^{-k\tau}(1 - k\tau)$$

The maximum occurs at

$$\tau = 1/k$$

Thus the repeated-root branch naturally generates a humped maturity loading.

This is the first appearance of the humped structure as an admissibility result. It has not been imposed as a volatility shape. It has appeared because the two admissible exponential modes have coalesced.

5.8 The r = 1 branch

Under the normalisation used in this monograph, the repeated-root case is the r = 1 branch.

This branch has special importance because it generates the linear-exponential loading

$$H(\tau) \propto \tau e^{-k\tau}$$

Under the parameter convention used later, the repeated-root decay rate will be written as

$$k = a^2,$$

so that the humped component becomes

$$H(\tau) \propto \tau e^{-a^2\tau}$$

At this stage, this is an interior admissibility statement, not yet a boundary-condition statement. The r = 1 branch identifies a humped mode that is allowed by the interior law. Whether that mode is realised depends on boundary closure.

This distinction is central:

interior admissibility produces the repeated-root mode;

boundary closure selects its economic realisation.

The r = 1 branch is therefore the bridge between the coefficient-matching algebra and the later boundary-law interpretation of humped volatility.

5.9 Relation to humped volatility

Humped volatility profiles are familiar in interest-rate modelling. In many models, such a hump is specified directly because it fits observed term-structure volatility.

The IRFR interpretation is different. The repeated-root mode shows that a humped maturity profile can arise from the admissibility structure of the field itself.

The key object is

$$h(\tau) = \tau e^{-k\tau}$$

This profile has three features:

$$h(0) = 0;$$

$$h(\tau) \text{ has an interior maximum;}$$

$$h(\tau) \to 0 \quad \text{as} \quad \tau \to \infty.$$

These are precisely the qualitative features associated with humped forward-rate volatility loadings: little or no instantaneous short-end volatility under a pinned boundary, maximal volatility at an intermediate maturity, and decaying sensitivity at long maturities.

This interpretation should be stated carefully. The existence of a humped admissible mode does not by itself prove that observed volatility surfaces must take this form. It shows that such a shape is structurally available within the IRFR field and can be selected by appropriate boundary closure.

Empirical relevance remains a calibration question. Structural admissibility is the mathematical result.

5.10 Interior modes and boundary selection

The analysis of this chapter reinforces the division introduced earlier.

The interior law determines the admissible maturity modes. These modes arise from the KFE, moment closure, and coefficient matching. They are mathematical consequences of the first three axioms.

The boundary law selects from these modes. It determines which admissible combination is realised at the short end and how the field is closed.

For the distinct-root branch, the boundary condition selects the constants in the two-exponential combination,

$$H(\tau) = C_1 e^{-k_1\tau} + C_2 e^{-k_2\tau}$$

For the repeated-root branch, the boundary condition selects the constants in

$$H(\tau) = (C_1 + C_2\tau)e^{-k\tau}$$

The Dirichlet condition, studied later, is especially important because it naturally suppresses the constant term and leaves the humped linear-exponential component.

Thus the boundary law does not merely impose an endpoint value. It selects the maturity structure of risk.

5.11 Relation to Ritchken–Chuang

This gives the repeated-root mode a natural connection with the humped-volatility form studied by Ritchken and Chuang [13, 14]. The connection is structural rather than model-equivalence: Ritchken and Chuang build a full option-pricing framework, while IRFR identifies a similar humped loading as an admissible boundary-selected mode. The point is therefore not that IRFR reproduces their pricing model, but that it supplies an endogenous field-theoretic route to a volatility shape that they treated within an interest-rate derivative-pricing setting.

This does not place the IRFR model outside HJM. The resulting dynamics should still admit an HJM representation under translation. The point is that the HJM representation is not structure-preserving: it does not reveal why the humped form arises.

The IRFR interpretation is therefore:

humped volatility as boundary-selected admissible mode;

rather than

humped volatility as exogenous volatility input.

The full monetary-policy interpretation will require the boundary-condition analysis developed later.

5.12 Summary

This chapter has extended the affine-closure analysis of Chapter 4 to the maturity loading.

The KFE and moment equations do not allow arbitrary maturity loadings. Coefficient matching imposes a second-order maturity equation. The solutions of that equation define the admissible modes of the IRFR field.

In the generic case, the characteristic roots are distinct, producing a two-exponential branch,

$$H(\tau) = C_1 e^{-k_1\tau} + C_2 e^{-k_2\tau}$$

At the repeated-root value, identified under the later normalisation with r = 1, the admissible loading becomes

$$H(\tau) = (C_1 + C_2\tau)e^{-k\tau}$$

The term

$$\tau e^{-k\tau}$$

is humped. It is therefore the structural origin of the humped volatility branch that later connects to Ritchken–Chuang-type specifications and to the Dirichlet boundary interpretation.

The main message is that humped maturity loading is not inserted by hand. It arises as a repeated-root admissible mode of the IRFR field.

5.13 Transition to boundary closure

The next chapter turns from interior admissibility to boundary selection.

The present chapter identified which maturity modes are allowed by the interior law. It did not yet determine which mode is realised. That selection requires Axiom 4: boundary closure at $\tau = 0$.

The next chapter therefore studies boundary conditions as economic regimes. It explains how natural, Dirichlet, and Robin closures select different admissible combinations of the maturity modes derived here, and why the Dirichlet repeated-root case plays a special role in the monetary-policy interpretation of humped volatility.

Chapter 6

Boundary Closure and Economic Regimes

6.1 Introduction

Chapter 5 derived the admissible maturity modes of the IRFR field. The interior law, generated by Axioms 1-3, restricts the maturity loading to a family of admissible modes. In the generic case, these modes are combinations of distinct exponentials. At the repeated-root point, identified with the r = 1 branch under the normalisation used in this monograph, the admissible loading contains the linear-exponential component

$$\tau \exp(-a^2\tau)$$

This humped component is an interior admissibility result. It is allowed by the dynamics, but it is not yet selected.

The present chapter turns to selection.

The IRFR field is defined on the half-line

$$\tau \in [0, \infty) \, ,$$

and all fixed-maturity points are transported toward the short end,

$$\tau = 0$$

Thus the field is not fully determined by its interior law. A boundary condition is required at the short end. This is the content of Axiom 4.

The key distinction is:

Axioms 1–3 determine admissible modes;

Axiom 4 selects the realised boundary regime.

Boundary closure is therefore not a technical afterthought. It is the mechanism by which economic regime enters the IRFR model.

6.2 Why boundary closure is necessary

The maturity domain of the IRFR field is the half-line. Unlike a model defined on the whole real line, a half-line field has an endpoint. That endpoint is not arbitrary. It is the instantaneous short end of the term structure.

The rolling-maturity identity

$$d\tau/dt = -1$$

means that every fixed calendar maturity moves toward this endpoint as calendar time advances. The short end is therefore the location at which the rolling field must be closed.

The interior dynamics derived in the previous chapters specify how the field evolves for positive remaining maturity. They determine the admissible forms of drift, diffusion, and maturity loading away from the boundary. But they do not, by themselves, determine what happens at $\tau = 0$.

This is a standard feature of boundary-value problems. An interior differential equation determines a space of possible solutions. A boundary condition selects the particular solution, or the particular family of solutions, compatible with the endpoint constraint.

In the IRFR setting, this mathematical fact has direct economic meaning. The short end is where policy, funding conditions, and overnight-rate mechanisms enter most directly. Boundary closure is therefore both a mathematical necessity and an economic modelling choice.

6.3 Interior admissibility versus boundary selection

The distinction between admissibility and selection is central.

Interior admissibility answers the question:

Which maturity loadings are dynamically allowed?

Boundary selection answers the question:

Which admissible loading is realised under a given short-end regime?

Chapter 5 answered the first question. It showed that admissible loadings arise as solutions of a second-order maturity equation. In the distinct-root case,

$$H(\tau) = C_1 \exp(-k_1\tau) + C_2 \exp(-k_2\tau)$$

In the repeated-root case,

$$H(\tau) = (C_1 + C_2\tau) \exp(-k\tau)$$

The constants C_1 and C_2 are not fixed by the interior equation alone. They encode the remaining freedom after admissibility has been imposed. Boundary closure fixes, constrains, or relates these constants.

Thus Axiom 4 does not create the admissible modes. It selects among them.

This point prevents a common misunderstanding. The humped repeated-root loading is not imposed by the boundary condition from nothing. It first appears as an admissible interior mode. The boundary condition determines whether that mode is selected, suppressed, or mixed with another admissible component.

6.4 The short end as an economic boundary

The short end of the yield curve has a special economic role.

Central banks operate most directly through short-maturity rates: overnight rates, reserve remuneration, standing facilities, policy corridors, and short-end guidance. These mechanisms do not usually set the entire forward curve directly. Instead, they act at or near the short end, and their effects propagate through the maturity structure.

In the IRFR formulation, this observation has a natural mathematical expression. Monetary policy enters through the boundary condition at

$$\tau = 0$$

This does not mean that every monetary-policy regime literally fixes the instantaneous short rate. Different policy regimes correspond to different forms of boundary closure. Some regimes leave the short end relatively free. Others pin it tightly. More realistic regimes partially anchor the short end while allowing residual stochastic variation.

The boundary condition is therefore the mathematical location of the policy regime.

This is the sense in which monetary policy is represented as a boundary law. It is not introduced as a global curve-wide force. It is introduced as a closure rule at the endpoint of the maturity-domain field.

6.5 Three canonical boundary regimes

The monograph focuses on three canonical boundary regimes:

Neumann/free closure;

Dirichlet fixed-boundary closure;

Robin partial-anchoring closure.

These regimes should be read as structural idealisations, not literal institutional descriptions. They are not claimed to exhaust all possible IRFR boundary laws. They provide the three canonical cases studied in this monograph: a non-pinned benchmark, a pinned boundary, and a partially anchored boundary.

Boundary regime	**Representative loading / condition**	**Economic interpretation**	**Mode effect**
Neumann/free	$H_N(\tau) = A_N \exp(-a_N^2\tau)$	non-pinned benchmark	retains boundary-active exponential loading

Boundary regime	Representative loading / condition	Economic interpretation	Mode effect
Dirichlet	$H(0) = 0$	pinned short end	suppresses C_1 in the repeated-root case
Robin	$\alpha H(0) + \beta H'(0) = \gamma$	partial anchoring	mixes admissible components

Neumann/free closure[1] represents the benchmark in which no explicit monetary-policy intervention is imposed through the boundary law. It should not be read as the absence of monetary institutions or policy influence in the economy. Rather, it is the case in which the short end remains stochastic and is closed without being externally pinned.

A Dirichlet boundary represents an idealised fixed-policy regime. The short end is pinned at a prescribed value. This is an exaggerated policy intervention, but it is mathematically important because it reveals how a fixed short-end boundary selects admissible modes.

A Robin boundary represents partial anchoring. It links the value of the loading and its derivative at the boundary. Economically, this corresponds to a short end influenced by policy but not perfectly fixed.

The rest of this chapter explains how these canonical closures act on the admissible mode space.

6.6 Neumann/free closure

For the loading-level analysis developed in this monograph, the Neumann/free benchmark is represented by the free exponential short-end stochastic mode

$$H_N(\tau) = A_N e^{-a_N^2 \tau}$$

This is the retained exponential loading for the non-pinned benchmark. The terminology follows the convention fixed in Chapter 6.

This loading satisfies

$$H_N(0) = A_N$$

Thus, unless $A_N = 0$, the short end retains stochastic loading. The boundary is closed, but it is not pinned.

This choice should be understood carefully. It is not a theorem that every possible Neumann/free IRFR closure must take this exact exponential form. A full state-space formulation could express Neumann/free closure through a flux, regularity, or positivity-preserving endpoint condition. Rather, the exponential free-boundary loading is the benchmark used here to define the non-pinned regime against which Dirichlet and Robin closure are compared.

[1] The term Neumann/free denotes the non-pinned boundary benchmark in which the boundary-active exponential loading is retained. The word "Neumann" is used in this boundary-regime sense only: it should not be read as imposing the classical homogeneous Neumann condition $H'(0) = 0$ on the stochastic loading. The operative condition in the present monograph is non-pinning rather than zero derivative.

Economically, this Neumann/free benchmark represents the case of no explicit monetary-policy intervention through the boundary law. It is not a claim that monetary policy is absent in an institutional sense. It means that the model imposes no additional policy-pinning condition at $\tau = 0$. The short end remains part of the stochastic field rather than being externally pinned.

This regime is therefore a non-pinned benchmark closure. It allows stochastic exposure to enter immediately at the short end, rather than suppressing the boundary-active component.

6.7 Dirichlet closure

A Dirichlet boundary condition fixes the value of the field or loading at the boundary. In schematic form,

$$H(0) = 0,$$

or, in terms of the short-end field,

$$x(t, 0) = \text{prescribed boundary value.}$$

The exact form depends on whether the boundary condition is imposed on the field level, the displacement from equilibrium, or the maturity loading. But the structural meaning is the same: the boundary value is fixed.

Economically, the Dirichlet case represents an idealised pinned short-end regime. It is not intended as a literal description of all monetary-policy operations. Rather, it is a limiting case in which the short end is treated as fixed by policy.

This case is especially important for the repeated-root branch.

Recall that the repeated-root admissible loading has the form

$$H(\tau) = (C_1 + C_2\tau) \exp(-k\tau)$$

At the boundary,

$$H(0) = C_1$$

A Dirichlet condition of the form

$$H(0) = 0$$

therefore sets

$$C_1 = 0$$

The realised loading becomes

$$H(\tau) = C_2\tau \exp(-k\tau)$$

Thus the Dirichlet boundary selects the humped linear-exponential component.

This is the first explicit example of boundary selection. The repeated-root mode is admissible by the interior law. The Dirichlet boundary selects the humped component by suppressing the constant exponential term.

6.8 The Dirichlet repeated-root branch

The Dirichlet repeated-root case is central to the monograph.

Under the r=1 repeated-root branch, the admissible loading is

$$H(\tau) = (C_1 + C_2\tau)\exp(-a^2\tau)$$

Imposing the Dirichlet condition

$$H(0) = 0$$

gives

$$C_1 = 0,$$

and hence

$$H(\tau) = C_2\tau\exp(-a^2\tau)$$

This loading is humped. It vanishes at the boundary, rises to an interior maximum, and decays at long maturities.

The maximum occurs at

$$\tau = 1/a^2$$

This is the structural source of the humped volatility branch.

In a conventional HJM representation, a humped volatility loading of this type may appear as an exogenous maturity function. In the IRFR formulation, it is selected by the combination of interior admissibility and boundary closure:

$r = 1$ *repeated-root admissibility;*

Dirichlet short-end boundary closure.

The humped profile is therefore not introduced by hand. It is a boundary-selected admissible mode.

6.9 Markovianity and representation

The Dirichlet repeated-root branch also clarifies the Markovianity claim made in Chapter 1.

The claim is not finite-dimensional Markovianity in the ordinary short-rate sense. It is Markovianity of the boundary-value field evolution.

This distinction matters. In standard finite-dimensional HJM coordinates, a stationary humped volatility profile of this type is often treated as non-Markovian. The hump appears as an externally specified maturity structure that cannot generally be represented by a low-dimensional Markov state without additional machinery.

In the IRFR formulation, the same humped structure appears differently. It is a boundary-selected mode of a rolling maturity field. The state is the field, the short end is the boundary, and the humped loading is generated by the boundary-value structure.

The representation matters.

What appears as non-Markovian structure in one coordinate system may become a local boundary condition in another. This is why the IRFR geometry is not merely notational. It changes which features of the model are primitive and which are derived.

6.10 Robin closure

A Robin boundary condition mixes the value of the field and its derivative at the boundary. In schematic form,

$$\alpha H(0) + \beta H'(0) = \gamma,$$

where α, β, and γ encode the boundary regime.

Economically, this represents partial anchoring. The short end is influenced by policy, but not perfectly pinned. There remains residual stochastic movement at the boundary.

The Robin case is therefore the natural bridge between the Neumann/free benchmark and the fixed boundary. It allows a continuum of regimes between a free short end and a fully pinned short end.

In the mode language, the Robin condition selects a mixture of admissible modes. For the repeated-root branch,

$$H(\tau) = (C_1 + C_2\tau)\exp(-k\tau),$$

the boundary value and derivative are

$$H(0) = C_1,$$

and

$$H'(0) = C_2 - kC_1$$

A Robin condition therefore imposes a linear relation between C_1 and C_2. Unlike the Dirichlet case, it need not suppress the constant term completely. Instead, it selects a mixed loading containing both exponential and humped components.

This is economically important. A realistic policy regime may partially anchor the short end without eliminating all short-end variation. Robin closure captures this intermediate regime.

6.11 Boundary law as policy transmission

The boundary law determines how short-end policy enters the admissible maturity structure.

In a conventional reduced-form model, policy effects may be introduced through shifts in the short rate, changes in factor dynamics, or changes in volatility specification. In the IRFR model, policy enters through the boundary condition.

This has two consequences.

First, policy is local at the point of intervention. It enters at the short-end boundary rather than as a global curve-wide forcing term.

Second, the effect of policy is transmitted through the admissible maturity modes of the field. The boundary does not directly prescribe the entire curve. It selects a mode or combination of modes, and the interior law propagates that selection across maturity.

This is the boundary-law interpretation of monetary policy.

The short end is where policy acts most directly. The maturity field is the structure through which that action is transmitted.

6.12 Relation to the r-branch

Boundary closure interacts with the r-branch developed in Chapter 5.

The r-branch determines the admissible interior mode structure. The boundary condition determines the realised combination.

For $r \neq 1$, the admissible loading is a combination of distinct exponential modes. Boundary closure selects the constants in that combination.

For r=1, the repeated-root loading is

$$H(\tau) = (C_1 + C_2\tau)\exp(-a^2\tau)$$

The Neumann/free benchmark retains the boundary-active exponential stochastic mode,

$$H_N(\tau) = A_N \exp(-a_N^2\tau)$$

The Dirichlet condition selects the humped component

$$H(\tau) \propto \tau \exp(-a^2\tau)$$

The Robin condition selects a mixture of the constant exponential and humped components.

Thus the r-branch determines the available mode structure, while the boundary law determines the economic realisation of that structure.

6.13 Summary

This chapter has introduced boundary closure as the selection mechanism for the IRFR field.

The interior law, derived from Axioms 1-3, determines the admissible maturity modes. Axiom 4 determines how the field is closed at the short end. This closure selects the realised combination of admissible modes.

The main boundary regimes are Neumann/free closure, Dirichlet fixed-boundary closure, and Robin partial-anchoring closure.

The Neumann/free benchmark retains the boundary-active exponential loading

$$H_N(\tau) = A_N \exp(-a_N^2 \tau)$$

The Dirichlet repeated-root case is especially important. Under r=1, the Dirichlet condition suppresses the constant exponential term and leaves the humped loading

$$H(\tau) \propto \tau \exp(-a^2 \tau)$$

The main message is:

interior admissibility creates the mode space;

boundary closure selects the realised mode.

6.14 Transition to detailed boundary cases

The next chapters study the boundary regimes in detail.

The Neumann/free case provides the non-pinned benchmark closure. The Dirichlet case provides the idealised pinned-policy limit and selects the repeated-root humped branch. The Robin case provides the more realistic partially anchored policy regime.

These boundary cases will show how the same admissible interior law can generate different observable maturity loadings depending on how the short end is closed.

Chapter 7

Neumann/free Boundary Closure and the Free Benchmark

7.1 Introduction

Chapter 6 introduced boundary closure as the selection mechanism for the IRFR field. The interior law determines the admissible maturity modes, while the boundary law determines how the short end is closed.

This chapter studies the benchmark case: Neumann/free boundary closure.

The Neumann/free boundary is the regime in which the short end is not externally pinned. The field is allowed to meet the boundary without imposing a fixed short-end value. In this sense, Neumann/free closure represents the non-pinned benchmark against which the Dirichlet and Robin regimes are compared.

Neumann/free closure leaves the short end stochastic;

Dirichlet closure pins the short end;

Robin closure partially anchors the short end.

The purpose of the present chapter is to establish the first case clearly. Neumann/free closure provides the baseline regime in which the short end remains active and stochastic. It does not select the pure humped repeated-root branch in the same way as the Dirichlet condition. Instead, it retains the ordinary boundary-active stochastic mode.

This benchmark is important for two reasons. First, it separates interior admissibility from policy selection. A humped repeated-root mode may be admissible, but it is not automatically realised. Its selection depends on how the short end is closed. Second, it clarifies what monetary policy changes. A pinned or partially anchored policy regime is meaningful only relative to a benchmark in which the short end remains free.

7.2 What "Neumann/free" means

For the loading-level analysis developed in this monograph, Neumann/free boundary closure is represented by the Neumann/free benchmark in which the short end remains part of the stochastic field.

$$H_N(\tau) = A_N e^{-a_N^2 \tau}$$

This loading satisfies

$$H_N(0) = A_N$$

Thus, unless $A_N = 0$, the short end retains stochastic loading. The boundary is closed, but it is not pinned.

This choice should be understood carefully. It is not a theorem that every possible Neumann/free IRFR closure must take this exact exponential form. A full state-space formulation could express Neumann/free closure through a flux, regularity, or positivity-preserving endpoint condition. Rather, the exponential free-boundary loading is the benchmark used here to define the non-pinned regime against which Dirichlet and Robin closure are compared.

Economically, this Neumann/free benchmark represents the case of no explicit monetary-policy intervention through the boundary law. It is not a claim that monetary policy is absent in an institutional sense. It means that the model imposes no additional policy-pinning condition at $\tau = 0$. The short end remains part of the stochastic field rather than being externally pinned.

7.3 The short end under Neumann/free closure

The short end $\tau = 0$ is economically distinguished, but Neumann/free closure does not treat it as fixed.

Instead, the short end remains part of the stochastic field. It may move with the same underlying Brownian source as the rest of the curve, subject to the admissibility restrictions of the IRFR model.

This matters because the short end is where the difference between free and policy regimes becomes visible. Under a Neumann/free boundary, the model does not impose a fixed policy value at the boundary.

$$x(t, 0) = \text{fixed policy value}.$$

is not imposed. Nor is the loading-level condition

$$H(0) = 0$$

imposed. The short end is closed, but not pinned. This distinction will be important in the next chapter. Dirichlet closure suppresses the boundary-active stochastic component. Neumann/free closure, by contrast, retains it.

7.4 Neumann/free closure and probability flux

One way to understand Neumann/free closure is through flux.

The KFE derived in Chapter 4 can be written in terms of probability flow over the state variable. Under the convention

$$dx = -f\,dt + \sqrt{2}\,g\,dW_t,$$

the KFE is

$$\frac{\partial p}{\partial t} = \partial(fp)/\partial x + \partial^2(g^2 p)/\partial x^2$$

This may be written as a conservation equation,

$$\frac{\partial p}{\partial t} = -\frac{\partial J}{\partial x},$$

where J is the probability flux in the state variable. With the above sign convention,

$$J = -fp - \partial(g^2 p)/\partial x$$

A Neumann/free endpoint condition is often associated with vanishing boundary flux. In the state variable, this means that probability does not leak out of the admissible state domain.

This flux is a flux in the state variable x, not along the maturity coordinate τ. The analogous maturity-domain idea is that the field meets the short-end boundary through an internally compatible endpoint condition rather than through an externally imposed boundary value.

The precise flux or endpoint condition depends on the state domain and on the positivity or admissibility constraints imposed on the IRFR density. For the present chapter, the important point is conceptual: Neumann/free closure is a compatibility condition, not a pinned-value condition. It allows the field to remain stochastic at the boundary while preserving admissibility.

7.5 Neumann/free closure and the free stochastic mode

The Neumann/free benchmark retains the ordinary exponential stochastic mode. In the r=1 benchmark, the loading is

$$H_N(\tau) = A_N e^{-a_N^2 \tau}$$

The associated instantaneous variance rate is

$$\partial\,\mathrm{Var}(x^N_{t,t+\tau})/\partial t = 2a_N^2 A_N^2 e^{-2a_N^2 \tau}$$

This profile is monotone in maturity and non-zero at the short end whenever $A_N \neq 0$. It therefore represents a regime in which short-end stochastic exposure is retained.

This is the important distinction from Dirichlet closure. Neumann/free closure does not suppress the boundary-active exponential component. It allows stochastic exposure to enter immediately at the short end.

7.6 Benchmark interpretation

The Neumann/free benchmark is not the pure humped repeated-root branch. Its loading is

$$H_N(\tau) = A_N e^{-a_N^2 \tau}$$

By contrast, the Dirichlet repeated-root loading is

$$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau}$$

The two regimes therefore differ sharply at the boundary:

$$H_N(0) = A_N,$$

$$H_D^{(1)}(0) = 0$$

Neumann/free closure retains short-end loading. Dirichlet closure suppresses it.

This makes the economic distinction transparent. Under Neumann/free closure, the short end remains stochastic. Under Dirichlet closure, the short-end stochastic loading is removed by the pinned boundary.

7.7 Neumann/free closure versus Dirichlet selection

The comparison with Dirichlet closure is instructive.

The Neumann/free benchmark has loading

$$H_N(\tau) = A_N e^{-a_N^2 \tau}$$

The Dirichlet repeated-root branch has loading

$$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau}$$

The former is boundary-active, while the latter is boundary-suppressed:

$H_N(0) = A_N, \quad H_D^{(1)}(0) = 0.$

Boundary regime	Loading	Short-end loading	Economic meaning
Neumann/free	$A_N e^{-a_N^2 \tau}$	retained	non-pinned benchmark

Boundary regime	Loading	Short-end loading	Economic meaning
Dirichlet	$C_D \tau e^{-a_D^2 \tau}$	suppressed	pinned boundary

This is the point of boundary law. The short-end regime changes the realised loading, and therefore changes how stochastic sensitivity is distributed across maturity.

The difference has economic meaning. Dirichlet closure removes short-end stochastic loading and represents an idealised pinned policy boundary. Neumann/free closure preserves short-end variation and represents the non-pinned benchmark.

7.8 Neumann/free closure and the benchmark economy

The Neumann/free boundary is best interpreted as a benchmark economy rather than as a literal absence of all policy.

No real interest-rate system is free of institutional structure. Overnight markets, collateral systems, reserve regimes, and central-bank operating frameworks always matter. Nevertheless, the Neumann/free boundary is useful because it describes the limiting case in which the short end is not pinned by an externally prescribed value.

It is the regime in which the field is closed by internal admissibility rather than by policy anchoring. This makes it the appropriate baseline for the monetary-policy extensions.

Neumann/free closure: free short end;

Dirichlet closure: pinned short end;

Robin closure: partially anchored short end.

The benchmark therefore defines what policy intervention changes. It changes the boundary law.

7.9 Consequences for volatility shape

Neumann/free closure does not imply a flat volatility structure. The Neumann/free benchmark has a non-trivial maturity profile, but it is the boundary-active exponential profile rather than the pinned-boundary hump.

$H_N(\tau) = A_N \, e^{-a_N^2 \tau}$.

The pure humped Dirichlet profile

$$\tau e^{-a^2 \tau}$$

is specifically associated with suppressing the boundary-active component and selecting the repeated-root boundary mode.

Neumann/free closure can produce decay and maturity dependence. But it does not express the same economic content as a pinned short end. The short-end loading remains present.

This matters for empirical interpretation. Observed volatility near the short end contains information about the boundary law, not merely about the interior dynamics.

7.10 Relation to monetary policy

The Neumann/free boundary provides the reference point for monetary-policy interpretation.

In a free-boundary benchmark, the short end remains stochastic. Policy is not represented as fixing the boundary value. The field evolves according to the interior admissibility structure and the Neumann/free closure condition.

When policy becomes active in the boundary-law sense, the closure changes. A Dirichlet boundary represents the limiting case in which the short end is pinned. A Robin boundary represents the more realistic case in which the short end is partially anchored.

Thus monetary policy is not modelled as a new interior force. It is modelled as a change in the boundary law.

interior dynamics describe admissible propagation;

boundary law describes regime selection.

The Neumann/free boundary is the non-pinned benchmark from which policy anchoring is measured.

7.11 Summary

This chapter has studied Neumann/free boundary closure as the benchmark regime of the IRFR field.

Neumann/free closure is required because the maturity domain is the half-line and the short end is a boundary. But unlike Dirichlet closure, it does not prescribe a fixed boundary value. It allows the short end to remain stochastic.

In the loading-level benchmark used here, Neumann/free closure retains the ordinary exponential short-end stochastic mode,

$H_N(\tau) = A_N\, e^{-a_N^2 \tau}$.

This contrasts with the Dirichlet repeated-root loading,

$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau}$,

which suppresses short-end loading.

Neumann/free closure closes the field without pinning the short end.

7.12 Transition to Dirichlet closure

The next chapter studies the Dirichlet boundary in detail.

Dirichlet closure fixes the boundary value at the short end. In the repeated-root r=1 branch, this suppresses the boundary-active stochastic component and selects the humped loading

$$H(\tau) \propto \tau e^{-a^2 \tau}$$

This case represents the idealised pinned-policy limit. It is mathematically simple, economically extreme, and central to the connection between IRFR boundary law and humped volatility.

Chapter 8

Dirichlet Boundary Closure and the r = 1 Humped Branch

8.1 Introduction

Chapter 7 studied Neumann/free boundary closure as the benchmark regime of the IRFR field. Neumann/free closure closes the maturity-domain field without pinning the short end. The short end remains stochastic, and the boundary-active exponential loading is retained.

This chapter studies the opposite limiting case: Dirichlet boundary closure.

A Dirichlet boundary fixes the value of the field, or of the relevant displacement or loading, at the short-end boundary. In schematic loading form,

$$H(0) = 0,$$

or, at the field level,

$$x(t, 0) = \text{prescribed boundary value.}$$

The exact object to which the condition is applied depends on the formulation. The structural meaning is the same: the boundary value is fixed.

Economically, the Dirichlet case represents an idealised pinned-policy regime. It is not intended as a literal description of all central-bank operations. It is a limiting case in which the short end is treated as fixed by policy. Precisely because it is extreme, it is analytically revealing.

The central result of the chapter is that, in the repeated-root r=1 branch, Dirichlet closure selects the humped linear-exponential loading

$$H(\tau) \propto \tau e^{-a^2\tau}$$

The humped loading is therefore not an exogenous volatility input. It is the boundary-selected survivor of the repeated-root admissible mode space.

8.2 Dirichlet closure as fixed-boundary selection

The IRFR maturity domain is the half-line

$$\tau \in [0, \infty) .$$

The endpoint $\tau = 0$ is the instantaneous short end. A Dirichlet condition fixes the value of the relevant boundary object at that endpoint.

In the simplest loading formulation, the condition is

$$H(0) = 0$$

This says that the maturity loading vanishes at the boundary. In economic terms, the short-end stochastic loading is suppressed. The boundary is not merely closed; it is pinned.

At the field level, a Dirichlet condition may instead be written as

$$x(t, 0) = x_0(t),$$

where $x_0(t)$ is a prescribed short-end boundary value. In the most idealised fixed-policy case, $x_0(t)$ may be constant over the relevant horizon. More generally, it may be a deterministic policy path.

In the loading-level analysis below, we use the corresponding homogeneous boundary condition H(0)=0, which represents the suppression of short-end stochastic loading.

For the mode-selection argument, the essential feature is that the boundary value is prescribed externally. This distinguishes Dirichlet closure from Neumann/free closure, which retains the boundary-active stochastic component.

8.3 Economic interpretation: the pinned short end

The Dirichlet case should be interpreted as an idealised pinned short-end regime.

Central banks act most directly at the short end. In a policy-corridor, administered-rate, or fixed-policy-rate regime, the shortest maturities may be strongly anchored relative to the rest of the curve. The Dirichlet condition abstracts this anchoring into a fixed boundary value.

This abstraction is deliberately sharp. It suppresses short-end stochastic loading and asks what maturity structure remains admissible in the interior. The boundary condition therefore acts as a selection rule on the admissible maturity modes.

This is why the Dirichlet case is central to the monetary-policy extension. It provides the cleanest mathematical representation of extreme short-end policy anchoring.

8.4 Action on the repeated-root mode space

The repeated-root branch has admissible loading

$$H(\tau) = (C_1 + C_2\tau)e^{-k\tau}$$

Under the normalisation used for the r=1 branch,

$$k = a^2,$$

so

$$H(\tau) = (C_1 + C_2\tau)e^{-a^2\tau}$$

The two independent components are

$e^{-a^2\tau}$

and

$$\tau e^{-a^2\tau}$$

The first component is non-zero at the boundary, while the second component vanishes at the boundary:

$e^{-a^2 \cdot 0} = 1$, $\tau e^{-a^2\tau}|_{\tau=0} = 0$.

Thus the two repeated-root components behave differently at the short end. The Dirichlet condition distinguishes precisely between them.

A Dirichlet condition of the form H(0)=0 suppresses the constant exponential component and leaves the linear-exponential component. This is the algebraic mechanism by which the Dirichlet boundary selects the humped branch.

8.5 The selection calculation

Start with the repeated-root admissible loading

$$H(\tau) = (C_1 + C_2\tau)e^{-a^2\tau}$$

Evaluating at the boundary gives

$$H(0) = C_1$$

The Dirichlet condition H(0)=0 therefore implies

$C_1 = 0$.

The realised loading is

$$H(\tau) = C_2\tau e^{-a^2\tau}$$

Up to normalisation,

$$H(\tau) \propto \tau e^{-a^2\tau}$$

The humped term was already present in the interior admissible mode space. The Dirichlet boundary did not create it from nothing. It selected it by eliminating the boundary-active constant component.

8.6 Shape of the selected loading

Ignoring the normalisation constant, define

$$h(\tau) = \tau e^{-a^2\tau}$$

Then $h(0) = 0$ and $h(\tau) \to 0$ as $\tau \to \infty$. The derivative is

$h'(\tau) = e^{-a^2\tau}(1 - a^2\tau)$.

The maximum occurs when

$1 - a^2\tau = 0$,

so

$$\tau_{\max} = 1/a^2$$

Thus the selected loading vanishes at the short end, reaches an interior maximum, and decays at long maturity. This is the qualitative structure of a humped forward-rate volatility loading.

The location of the maximum is controlled by a^2. A larger a^2 shifts the hump toward shorter maturities. A smaller a^2 shifts it toward longer maturities. Thus a^2 controls not only the strength of local mean reversion in the affine drift operator, but also the maturity scale of the selected Dirichlet mode.

8.7 Positivity and the lower admissible boundary

The Dirichlet repeated-root branch also has an important admissibility property. The selected loading

$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2\tau}$

is non-negative on the half-line whenever $C_D \geq 0$. It satisfies

$$H_D^{(1)}(0) = 0, H_D^{(1)}(\tau) \geq 0 for \tau \geq 0, H_D^{(1)}(\tau) \to 0 as \tau \to \infty$$

Thus the boundary-selected loading does not force a sign change across maturity. In the Dirichlet regime, the short-end boundary can therefore be interpreted not only as a pinned policy value, but also as a lower admissible boundary for the rolling field.

Write the field relative to a minimum admissible level:

$$x(t,\tau) = x_{\min} + y(t,\tau), y(t,\tau) \geq 0$$

The Dirichlet condition is then imposed on the displacement,

$$y(t,0) = 0$$

For each fixed maturity $\tau > 0$, the shifted Dirichlet dynamics remain in the positive-state multiplicative KFE class considered in Appendix B. In schematic form,

$$dy_t = \kappa_D(\tau)(\theta_D(\tau) - y_t)dt + \sqrt{2}\eta_D(\tau)y_t\, dW_t,$$

with positive parameter functions. The stationary KFE then admits a shifted inverse-gamma density on y>0, equivalently on x>$x_{\min}$. The same Feller boundary classification as in the positive-state benchmark shows that the lower boundary y=0 is inaccessible for each fixed $\tau > 0$, provided the process starts in the positive state domain and the parameter conditions are satisfied.

The lower-boundary interpretation is therefore stronger than a static sign condition on the loading. It is supported by the shifted inverse-gamma stationary density and by maturity-by-maturity Feller inaccessibility of the lower state boundary. The detailed density and boundary-classification calculation is given in Appendix B.

The remaining qualification is field-level regularity. To pass from fixed-τ positivity to the full maturity-field statement

$$y(t,\tau) \geq 0 forall \tau \geq 0,$$

one also requires sufficient regularity of the field in the maturity coordinate. Under such regularity, the Dirichlet boundary gives y(t,0)=0 while the shifted process remains positive for each fixed $\tau > 0$.

This should be read as an admissibility statement about the IRFR field under the chosen boundary normalisation. It is not a claim that nominal market rates can never be negative in every institutional setting.

8.8 Comparison with Neumann/free closure

The contrast with Neumann/free closure is immediate.

The Neumann/free benchmark retains the boundary-active exponential stochastic mode:

$H_N(\tau) = A_N\, e^{-a_N^2 \tau}$.

The short-end loading remains present:

$H_N(0) = A_N$.

Under Dirichlet closure, the selected repeated-root loading is

$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau}$.

The short-end loading is suppressed:

$H_D^{(1)}(0) = 0.$

The difference is not merely a difference in shape. It is a difference in boundary selection.

Boundary regime	Loading	Short-end loading	Economic meaning
Neumann/free	$A_N e^{-a_N^2 \tau}$	retained	non-pinned benchmark
Dirichlet	$C_D \tau e^{-a_D^2 \tau}$	suppressed	pinned boundary

This is why the Dirichlet boundary has a natural interpretation as a pinned-policy limit.

8.9 Markovianity of the Dirichlet humped branch

The Dirichlet repeated-root branch clarifies the Markovianity issue.

In a standard HJM representation, a stationary humped volatility loading of the form

$$\tau e^{-a^2 \tau}$$

may appear as an exogenous maturity-dependent volatility specification. In finite-dimensional HJM realisation theory, such a specification is often treated as difficult to represent by a low-dimensional Markov state.

The IRFR interpretation is different. In IRFR coordinates, the humped loading is not an arbitrary maturity function. It is a boundary-selected admissible mode of a rolling maturity field. The state is the field, the short end is the boundary, and the boundary condition is local.

Thus the claim is not that the model is one-dimensional Markov in the ordinary short-rate sense. It is not. The claim is that the Dirichlet humped branch is Markovian as a boundary-value evolution of the IRFR field.

The Markovianity claim should be understood in the field-state sense. The state at time t is not a single short rate, nor a finite vector of HJM factors, but the rolling field

$X_t = \{x(t, \tau) : \tau \geq 0\},$

together with the boundary regime defining the domain of admissible evolution. With respect to the filtration generated by the driving Brownian motion and the current field state, the future law of X_{t+s} is determined by X_t, the interior operator, and the boundary condition. In this sense the Dirichlet repeated-root branch is Markovian as an infinite-dimensional boundary-value field evolution. This is weaker than finite-dimensional Markovianity, and it is the only Markovianity claim being made here.

8.10 Relation to Ritchken–Chuang-type volatility

The selected Dirichlet loading

$$H(\tau) \propto \tau e^{-a^2 \tau}$$

has the same qualitative form as the humped maturity loading associated with Ritchken–Chuang-type volatility specifications [13, 14].

In the conventional use of such specifications, the humped form is introduced directly as a maturity function. The IRFR interpretation is different. The same qualitative form appears as the result of boundary selection: r=1 repeated-root admissibility together with Dirichlet boundary closure.

This does not place the IRFR model outside HJM. The resulting dynamics should still admit an HJM representation after translation into standard forward-rate coordinates. The distinction is that the HJM representation does not reveal why the humped form arises.

In IRFR coordinates, the humped volatility profile is not an exogenous input. It is a boundary-selected admissible mode.

8.11 Why the Dirichlet case is extreme

The Dirichlet boundary is analytically clean, but economically extreme.

A real central bank rarely fixes the short end in a perfectly deterministic way over all relevant horizons. Even in strongly anchored regimes, there may be residual uncertainty about future policy, liquidity conditions, corridor implementation, and market microstructure.

The Dirichlet case should therefore be interpreted as a limiting regime. It identifies what happens when short-end stochastic loading is fully suppressed.

This extreme case is valuable because it reveals the mechanism clearly. Once the short-end loading is forced to vanish, the repeated-root branch leaves the humped component as the surviving admissible mode.

More realistic regimes will generally be mixed. They will allow some short-end variation while still anchoring the boundary. This is the role of Robin closure, studied later.

8.12 Boundary law and policy transmission

The Dirichlet case provides the sharpest illustration of policy as boundary law.

The policy action is local: it pins the short-end boundary. The consequence is non-local: the selected admissible mode determines how stochastic sensitivity is distributed across the maturity domain.

The boundary condition does not directly prescribe the entire curve. It selects a mode. The interior law then propagates that selection across maturity.

policy enters locally at $\tau = 0$;

the admissible field transmits the effect across $\tau > 0$.

In the Dirichlet repeated-root case, this transmission takes the form of a humped maturity loading. The short end is suppressed, intermediate maturities carry the greatest sensitivity, and long maturities decay.

8.13 Summary

This chapter has studied Dirichlet boundary closure as the idealised pinned short-end regime.

The repeated-root r=1 admissible loading is

$$H(\tau) = (C_1 + C_2\tau)e^{-a^2\tau}$$

Dirichlet closure imposes

$$H(0) = 0$$

Since H(0)=C_1, the boundary condition gives $C_1 = 0$. The selected loading is therefore

$$H(\tau) = C_2\tau e^{-a^2\tau}$$

This loading is humped, with maximum at

$$\tau = 1/a^2$$

The Dirichlet r=1 branch also admits a lower-boundary interpretation. In shifted coordinates, the branch supports a shifted inverse-gamma stationary density for fixed maturities $\tau > 0$, and the lower boundary is inaccessible maturity-by-maturity under the Feller conditions. Full field-level lower-boundary preservation requires the corresponding maturity-regularity assumption.

Dirichlet closure selects the humped repeated-root mode.

8.14 Transition to Robin closure

The next chapter studies Robin boundary closure.

Robin closure relaxes the extreme assumption of a fully pinned short end. Instead of fixing the boundary value outright, it allows a residual boundary-active component while adding a boundary-induced repeated-root component:

$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2\tau}$.

Economically, this represents partial anchoring. The short end is influenced by policy but not perfectly fixed. In the repeated-root branch, Robin closure therefore provides the more realistic bridge between the Neumann/free benchmark and the idealised Dirichlet limit.

Chapter 9

Robin Boundary Closure and Partial Policy Anchoring

9.1 Introduction

Chapter 8 studied Dirichlet boundary closure as the idealised pinned-policy limit. In the repeated-root r=1 branch, the Dirichlet condition suppresses short-end stochastic loading and selects the humped loading

$$H(\tau) \propto \tau e^{-a^2\tau}$$

This is the cleanest example of boundary selection: the interior equation creates the repeated-root mode space, while the boundary condition selects the component compatible with a pinned short end.

The Dirichlet case is analytically revealing, but economically extreme. A real policy regime rarely fixes the short end perfectly. Even when the policy rate is strongly anchored, the market short end may still reflect liquidity conditions, corridor mechanics, expectations of future policy, balance-sheet constraints, and residual stochastic variation.

This chapter studies the intermediate case: Robin boundary closure.

Robin closure retains residual short-end stochastic loading while allowing a boundary-induced repeated-root component. Algebraically, in the repeated-root branch this gives a mixed loading of the form

$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2\tau}$.

Economically, Robin closure represents partial policy anchoring. The short end is influenced by policy but not completely fixed.

The repeated-root r=1 branch is not studied here because empirical humped volatility proves that this branch is the correct one. Rather, it is studied because it is a mathematically distinguished branch of the admissibility equation: the two exponential modes coalesce, producing the linear-exponential component. The fact that this component resembles humped volatility forms used in interest-rate modelling makes the branch economically interpretable and worthy of detailed study. It does not, by itself, establish that r=1 is empirically preferred.

The role of the Robin analysis is therefore conditional and precise. It shows that, within the r=1 repeated-root mode space, partial boundary anchoring selects a family of mixtures between a boundary-active exponential component and a humped linear-exponential component.

9.2 Robin closure as a mixed boundary condition

A Robin boundary condition imposes a linear relation between the value of the loading and its derivative at the boundary. In homogeneous form,

$\alpha H(0) + \beta H'(0) = 0.$

More generally,

$\alpha H(0) + \beta H'(0) = \gamma.$

Economically, the Robin case represents partial anchoring. The short end is not fully pinned, but neither is it left entirely unconstrained by the boundary law.

It is useful to distinguish two ideas. First, the Neumann/free benchmark retains the ordinary exponential short-end stochastic mode,

$H_N(\tau) = A_N\, e^{-a_N^2\tau}.$

Second, the Dirichlet repeated-root boundary suppresses short-end loading and selects

$H_D^{(1)}(\tau) = C_D\tau e^{-a_D^2\tau}.$

Robin closure lies between these cases economically. It retains residual short-end loading while also allowing a boundary-induced repeated-root component. Algebraically, this gives

$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2\tau}.$

Thus Robin closure is best understood as the partially anchored repeated-root regime. It does not simply equal the Neumann/free benchmark, and it does not fully collapse to the Dirichlet branch unless the residual short-end component is suppressed.

9.3 Economic interpretation: partial anchoring

The economic interpretation of Robin closure is partial policy anchoring.

Under Neumann/free closure, the short end remains part of the stochastic field. It is closed, but not pinned. Under Dirichlet closure, the short end is pinned and short-end stochastic loading is suppressed.

Robin closure lies between these regimes. It allows the short end to be influenced by policy while retaining residual stochastic movement.

This is closer to how many real monetary-policy regimes operate. A central bank may strongly influence overnight or short-maturity rates, but the market short end may still move because of liquidity premia, expectations, collateral conditions, implementation frictions, and uncertainty about future policy.

The Robin condition captures this by imposing not a fixed value, but a relation between the boundary level and the boundary slope. Economically, this says that the short end is anchored relative to the way the curve leaves the boundary.

the boundary is constrained, but not fixed.

9.4 Action on the repeated-root mode space

The repeated-root Robin loading is

$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2\tau}$.

The two components are

$C_0 e^{-a_R^2\tau}$

and

$C_1\tau e^{-a_R^2\tau}$.

The first is boundary-active:

$H_R(0) = C_0$.

The second is boundary-induced and vanishes at the short end.

Differentiating,

$H_R'(\tau) = [C_1 - a_R^2(C_0 + C_1\tau)]e^{-a_R^2\tau}$,

so

$H_R'(0) = C_1 - a_R^2 C_0$.

A homogeneous Robin condition may be written as

$H_R'(0) = -\rho H_R(0)$,

where

$\rho = -H_R'(0)/H_R(0)$

whenever $H_R(0) \neq 0$. Substituting the boundary value and derivative gives

$$\rho = a_R^2 - C_1/C_0$$

Equivalently,

$$C_1/C_0 = a_R^2 - \rho$$

Thus the Robin loading can be written as

$H_R(\tau) = C_0[1 + (a_R^2 - \rho)\tau]e^{-a_R^2\tau}$.

Writing k=a_R^2, this is

$H_R(\tau) = C_0[1 + (k - \rho)\tau]e^{-k\tau}$.

The parameter ρ measures the negative boundary slope of the loading relative to its short-end level. It therefore controls the mixture of boundary-active exponential loading and boundary-induced humped loading.

9.5 Limiting cases and reference regimes

The Robin repeated-root loading is

$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2\tau}$.

This form contains two components:

$e^{-a_R^2\tau}$, $\tau e^{-a_R^2\tau}$.

If $C_1 = 0$, the loading is purely exponential:

$H_R(\tau) = C_0 e^{-a_R^2\tau}$.

This is the boundary-active profile associated with retained short-end stochastic loading. It aligns with the Neumann/free benchmark.

If $C_0 = 0$, the loading becomes

$H_R(\tau) = C_1\tau e^{-a_R^2\tau}$.

This is the Dirichlet repeated-root humped loading.

Thus Robin closure should be read as a mixed repeated-root family. It interpolates economically between retained short-end stochastic loading and pinned-boundary suppression by varying the relative weights of C_0 and C_1.

This is an economic interpolation, not a claim that all three cases are obtained as finite parameter values of one homogeneous Robin formula.

Regime	**Loading**	**Short-end loading**	**Interpretation**
Neumann/free	$A_N e^{-a_N^2\tau}$	retained	non-pinned benchmark
Robin	$(C_0 + C_1\tau)e^{-a_R^2\tau}$	partially retained	partial anchoring
Dirichlet	$C_D\tau e^{-a_D^2\tau}$	suppressed	pinned boundary

9.6 The anchoring parameter

The Robin anchoring parameter is

$\rho = -H_R'(0)/H_R(0)$.

For

$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2\tau}$,

we have

$H_R(0) = C_0$, $H_R'(0) = C_1 - a_R^2 C_0$.

Therefore,

$\rho = a_R^2 - C_1/C_0$.

Equivalently,

C_1=(a_R^2-ρ)C_0.

This gives

$H_R(\tau) = C_0[1 + (a_R^2 - \rho)\tau]e^{-a_R^2\tau}$.

The parameter ρ controls how the loading leaves the boundary. If ρ=a_R^2, then $C_1 = 0$, and the loading is pure exponential. If ρ<a_R^2, then C_1>0, and the loading contains a positive boundary-induced linear-exponential component. If ρ>a_R^2, then C_1<0, and the loading may eventually change sign.

9.7 Shape of the Robin-selected loading

Write

$$k = a_R^2, q = k - \rho$$

Then

$H_R(\tau) = C_0(1 + q\tau)e^{-k\tau}$.

The derivative is

$H_R'(\tau) = C_0[(q - k) - kq\tau]e^{-k\tau}$.

An interior stationary point occurs when

(q-k)-kq$\tau = 0$.

If q$\neq$0, this gives

$$\tau_* = (q - k)/(kq) = -\rho/[k(k - \rho)]$$

The shape depends on the anchoring parameter. If ρ=a_R^2, then q=0 and the loading is pure exponential. If ρ<a_R^2, then q>0 and the loading remains globally non-negative for $C_0 > 0$. If ρ>a_R^2, then q<0 and the loading eventually changes sign.

The Robin family can therefore produce boundary-active decay, shoulder-like profiles, or partially humped profiles. Unlike the Dirichlet branch, it need not vanish at the short end.

In the pure Dirichlet branch,

$H_D^{(1)}(\tau) \propto \tau e^{-k\tau}$,

and the maximum is necessarily at

$$\tau = 1/k$$

In the Robin case, the curve may retain short-end loading and may or may not display a pronounced interior hump, depending on the mixture selected by the boundary condition.

9.8 Positivity and admissibility under Robin closure

The Dirichlet r=1 branch has the cleanest lower-boundary interpretation. Its selected loading

$H_D^{(1)}(\tau) = C_D \tau e^{-k\tau}$

is non-negative on the half-line for $C_D \geq 0$ and vanishes at the boundary. In the shifted positive-state formulation, Appendix B shows that the Dirichlet branch admits a shifted inverse-gamma stationary density for each fixed $\tau > 0$ and that the lower boundary is inaccessible maturity-by-maturity under the Feller conditions.

Robin closure is more delicate because it retains residual short-end loading. The selected loading

$H_R(\tau) = C_0(1 + q\tau)e^{-k\tau}$

is non-negative on the half-line if the prefactor and the linear term preserve sign. If $C_0 > 0$, the necessary and sufficient condition for non-negativity on the entire half-line is

$q \geq 0.$

Equivalently,

$k - \rho \geq 0,\ \rho \leq k.$

If $\rho > k$, the linear term eventually becomes negative, and the loading changes sign at sufficiently large maturity. Such a branch may still have a mathematical interpretation, but it is not compatible with a simple non-negative lower-boundary interpretation on the whole half-line.

When the Robin loading is sign-compatible and the shifted positive-state KFE structure is retained, Robin inherits the same maturity-by-maturity inverse-gamma and Feller logic as the Dirichlet case. A separate density derivation is not required. The additional Robin requirement is the global sign condition on the loading.

Thus Robin admissibility has two layers: the local shifted-state diffusion must remain in the positive-state Feller class, and the maturity loading must remain sign-compatible on the domain under consideration. Full field-level lower-boundary preservation again requires the corresponding maturity-regularity assumption.

9.9 Robin closure and monetary-policy regimes

Robin closure gives a natural language for partially anchored monetary-policy regimes.

In a fully pinned regime, the short end is treated as fixed. This is the Dirichlet idealisation. In a free-boundary benchmark, the short end remains stochastic. This is Neumann/free closure. Between these extremes lies the more realistic case in which policy anchors the short end imperfectly.

The Robin condition captures this intermediate regime. The boundary value $H_R(0) = C_0$ represents the amount of short-end loading retained. The derivative $H_R'(0)$ represents how the maturity loading leaves the boundary. The Robin relation constrains these two quantities relative to one another.

Economically, this means that policy does not eliminate short-end stochastic variation entirely. Instead, it controls how short-end variation is transmitted into the rest of the curve.

The anchoring parameter ρ therefore acts as a transmission coefficient. It determines how much short-end stochastic variation enters the interior through the admissible mode mixture.

This is exactly the role of a partially anchored policy regime. Policy influences the boundary, but the market still supplies residual stochastic movement and slope at the short end.

9.10 Comparison with Neumann/free and Dirichlet regimes

The three canonical closures now have clearer interpretations.

Neumann/free closure retains the ordinary exponential short-end stochastic mode:

$H_N(\tau) = A_N\, e^{-a_N^2 \tau}$.

Dirichlet closure suppresses short-end loading and selects the repeated-root humped mode:

$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau}$.

Robin closure lies between these regimes by retaining a residual short-end component while adding a boundary-induced repeated-root component:

$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2 \tau}$.

Neumann/free $\to$ short-end stochastic loading retained;

Dirichlet $\to$ short-end stochastic loading suppressed;

Robin $\to$ short-end stochastic loading partially retained.

The same boundary-field logic therefore produces different observable maturity loadings depending on how the short end is closed.

9.11 Relation to humped volatility

Robin closure modifies the interpretation of humped volatility.

In the Dirichlet case, the humped loading is selected cleanly:

$H_D^{(1)}(\tau) \propto \tau e^{-k\tau}$.

In the Robin case, the selected loading is generally mixed:

$H_R(\tau) = (C_0 + C_1\tau)e^{-k\tau}$.

This may display hump-like behaviour, depending on the anchoring parameter and coefficient signs. But the hump is no longer forced to vanish at the short end. Instead, the profile may include a residual short-end component.

This is economically plausible. In partially anchored regimes, short maturities may retain some volatility, while intermediate maturities may still display heightened sensitivity. Robin closure therefore provides a structural route to shifted or partially humped volatility profiles.

This observation should be read carefully. The resemblance between Robin-selected profiles and empirical humped-volatility specifications motivates the study of the r=1 branch, but it does not validate that branch by itself. Empirical relevance remains a calibration and testing question.

Dirichlet: pure pinned-boundary hump;

Robin: mixed anchored-boundary profile.

9.12 Boundary law as a regime parameter

The Robin boundary condition suggests a broader modelling interpretation.

The boundary law itself can be regarded as a regime parameter. Changing the boundary coefficients, or equivalently changing the anchoring ratio ρ, changes the realised maturity loading without changing the interior admissible equation.

This is significant. It means that policy regimes can be modelled as changes in boundary law rather than changes in the interior dynamics.

The interior law describes the admissible propagation of stochastic shocks through maturity. The boundary law determines how the short end is closed and therefore which admissible profile is realised.

In empirical implementation, ρ becomes an estimable parameter linking observed maturity loadings to implicit policy anchoring strength. Different monetary-policy environments may therefore correspond not to different interior dynamics, but to different boundary-law selections within the same admissible branch.

observed maturity profile = interior admissible structure + boundary-law selection via ρ.

Thus changes in observed term-structure shape need not reflect changes in the underlying interior dynamics. They may instead reflect changes in the boundary regime.

9.13 Summary

This chapter has studied Robin boundary closure as the partially anchored regime of the IRFR field.

The repeated-root Robin loading is

$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2\tau}$.

The boundary value and slope are

$H_R(0) = C_0$, $H_R'(0) = C_1 - a_R^2 C_0$.

The anchoring parameter satisfies

$$\rho = a_R^2 - C_1/C_0,$$

or equivalently

$H_R(\tau) = C_0[1 + (a_R^2 - \rho)\tau]e^{-a_R^2\tau}$.

This loading retains short-end stochastic loading unless $C_0 = 0$. It therefore represents partial policy anchoring.

For $\rho \le a_R^2$, the selected loading can preserve a non-negative lower-boundary interpretation on the half-line, assuming $C_0 > 0$. Under the shifted positive-state KFE structure, the Robin case then inherits the fixed-maturity inverse-gamma and Feller logic described for Dirichlet in Appendix B. For $\rho > a_R^2$, the loading changes sign at sufficiently long maturity and requires a different interpretation.

Robin closure selects a mixed repeated-root loading.

This mixed loading lies between the Neumann/free benchmark and the Dirichlet pinned-boundary limit. It provides a structural representation of policy regimes in which the short end is influenced by policy but not perfectly fixed.

9.14 Transition

The next step is to connect the boundary-law taxonomy to observable term-structure behaviour.

The Neumann/free, Dirichlet, and Robin cases show that the same broad boundary-field logic can generate different maturity loadings depending on how the short end is closed. The next part of the monograph relates these boundary-selected loadings to observable volatility profiles, Ritchken–Chuang-type specifications, and empirical policy regimes.

The purpose is no longer only to derive admissible modes, but to ask how those modes can be identified, calibrated, and tested in market data.

Chapter 10

Boundary-Selected Loadings and Observable Volatility

10.1 Introduction

The previous chapters developed the boundary-law structure of the IRFR field.

Chapter 7 studied Neumann/free boundary closure as the non-pinned benchmark. Chapter 8 studied Dirichlet closure as the idealised pinned-policy limit. Chapter 9 studied Robin closure as the partially anchored regime.

Together, these chapters established the central boundary-law principle:

the interior law determines the admissible mode space;

the boundary law selects the realised maturity loading.

The present chapter translates that principle into observable term-structure behaviour. The object of interest is the maturity loading $H(\tau)$, which determines how the common stochastic shock is expressed across the curve. Under the one-factor IRFR convention,

$$g(\tau) = aH(\tau),$$

and the local stochastic evolution is written as

$$dx(t,\tau) = -f(x,t,\tau)dt + \sqrt{2}g(x,t,\tau)dW_t$$

When $g(\tau) = \mathrm{a}H(\tau)$, the instantaneous covariance between maturity points is

$$\mathrm{Cov}(dx(t,\tau_1), dx(t,\tau_2)) = 2a^2 H(\tau_1)H(\tau_2)dt$$

Thus $H(\tau)$ is the deterministic maturity profile through which the one stochastic factor is transmitted across the curve.

The purpose of this chapter is to interpret the boundary-selected forms of $H(\tau)$ as observable volatility loadings.

observable volatility shape is not imposed directly;

it is selected by boundary closure from the admissible mode space.

Equivalently, the volatility loading is the observable footprint of the boundary law.

10.2 From diffusion scale to volatility loading

In the IRFR convention, the diffusion-scale coefficient is g, while the conventional SDE volatility coefficient is $\sqrt{2}$g.

For the one-factor maturity-loading representation,

$$g(\tau) = aH(\tau)$$

Therefore the conventional instantaneous volatility loading is

$$\sqrt{2}aH(\tau)$$

The economically relevant maturity shape is $H(\tau)$. The constants a and $\sqrt{2}$ set scale and normalisation. The function $H(\tau)$ determines where along the curve stochastic sensitivity is concentrated.

This distinction matters. A one-factor model has only one Brownian source, but it can still generate non-trivial volatility shapes because the deterministic maturity loading can vary across τ.

Thus stochastic dimension and observed maturity profile are different objects. The covariance-rate kernel makes this explicit:

$$K(\tau_1, \tau_2) = (1/dt)\,\mathrm{Cov}(dx(t, \tau_1), dx(t, \tau_2)) = 2a^2 H(\tau_1)H(\tau_2)$$

The kernel is rank one, but its maturity dependence is shaped by H. The curve can therefore display structured maturity sensitivity even in a one-factor stochastic model.

In empirical covariance terms, the one-factor IRFR kernel has a single non-zero volatility direction. The corresponding maturity eigenfunction is proportional to $H(\tau)$. Thus, within the one-factor restriction, the boundary-selected loading may be compared directly with the leading empirical volatility mode of forward-rate changes.

10.3 Boundary-selected loading families

The boundary chapters produced three canonical loading families. They should not be presented as three finite parameter settings of a single homogeneous Robin formula. Rather, they are the observable loading forms associated with the three boundary regimes developed in Chapters 7-9.

The Neumann/free benchmark retains the ordinary exponential short-end stochastic loading:

$$H_N(\tau) = A_N \exp(-a_N^2 \tau)$$

The short-end loading is retained:

$$H_N(0) = A_N$$

Dirichlet closure suppresses short-end stochastic loading. In the repeated-root branch it selects the boundary-induced humped loading:

$$H_D^{(1)}(\tau) = C_D \tau \exp(-a_D^2 \tau)$$

The short-end loading is suppressed:

$$H_D^{(1)}(0) = 0$$

Robin closure represents partial anchoring. It retains a residual short-end component while allowing a boundary-induced repeated-root component:

$$H_R(\tau) = (C_0 + C_1 \tau) \exp(-a_R^2 \tau)$$

When $C_0 \neq 0$, this can also be written in anchoring-ratio form as

$$H_R(\tau) = C_0[1 + (a_R^2 - \rho)\tau] \exp(-a_R^2 \tau),$$

where

$$\rho = -H_R'(0)/H_R(0)$$

The three regimes may therefore be summarised as follows.

Boundary regime	**Loading**	**Short-end loading**	**Interpretation**
Neumann/free	$A_N \exp(-a_N^2 \tau)$	retained	non-pinned benchmark
Dirichlet	$C_D \tau \exp(-a_D^2 \tau)$	suppressed	pinned boundary
Robin	$(C_0 + C_1 \tau) \exp(-a_R^2 \tau)$	partially retained	partial anchoring

These forms are not unrelated volatility parametrisations. They are boundary-selected realisations of the IRFR maturity-domain structure.

10.4 The Neumann/free exponential benchmark

Neumann/free closure provides the non-pinned benchmark. It retains the boundary-active exponential mode

$$H_N(\tau) = A_N \exp(-a_N^2 \tau)$$

This loading is non-zero at the short end whenever $A_N \neq 0$. Thus short-end stochastic exposure is retained.

The associated instantaneous variance rate is

$$\partial_t \operatorname{Var}(x^N_{t,t+\tau}) = 2a_N^2 A_N^2 \exp(-2a_N^2\tau)$$

This profile is monotone in maturity. Its role is not to generate the pinned-boundary hump. Its role is to provide the free benchmark against which policy anchoring is measured.

Economically, the Neumann/free benchmark means that no additional policy-pinning condition is imposed at $\tau = 0$. It is not a claim that monetary policy is absent institutionally. It means that the boundary law does not suppress the short-end stochastic loading.

10.5 The Dirichlet humped loading

The Dirichlet case gives the cleanest humped volatility profile.

The selected repeated-root loading is

$$H_D^{(1)}(\tau) = C_D\tau \exp(-a_D^2\tau)$$

Ignoring the scale constant, define

$$h_D(\tau) = \tau \exp(-a_D^2\tau)$$

Then

$$h_D(0) = 0, h_D(\tau) \to 0 as \tau \to \infty$$

The derivative is

$$h_D'(\tau) = \exp(-a_D^2\tau)(1 - a_D^2\tau)$$

Thus

$$h_D'(\tau) = 0 \Leftrightarrow \tau = 1/a_D^2$$

The maximum occurs at

$$\tau_{\max} = 1/a_D^2$$

This is the standard qualitative humped shape: low short-end volatility, increasing sensitivity at intermediate maturities, and decay at long maturities.

The interpretive distinction is essential. In the IRFR framework, this shape is not specified as an empirical convenience. It is the boundary-selected survivor of the repeated-root mode space under a pinned short-end condition.

This is why the Dirichlet branch provides the cleanest structural route to humped volatility specifications.

10.6 The Robin mixed loading

The Robin case generalises the Dirichlet profile by retaining some boundary-active component.

The Robin-selected repeated-root loading is

$$H_R(\tau) = (C_0 + C_1\tau)\exp(-a_R^2\tau)$$

The two components are

$C_0 \exp(-a_R^2\tau)$

and

$C_1\tau \exp(-a_R^2\tau)$.

The first is boundary-active. The second is boundary-induced and vanishes at the short end.

At the boundary,

$$H_R(0) = C_0$$

Differentiating gives

$$H_R'(0) = C_1 - a_R^2 C_0$$

Thus, whenever $C_0 \neq 0$,

$$\rho = -H_R'(0)/H_R(0) = a_R^2 - C_1/C_0$$

Equivalently,

$$C_1/C_0 = a_R^2 - \rho$$

Therefore the Robin loading can be written as

$$H_R(\tau) = C_0[1 + (a_R^2 - \rho)\tau]\exp(-a_R^2\tau)$$

Writing k=a_R^2, this becomes

$H_R(\tau) = C_0[1 + (k - \rho)\tau]e^{-k\tau}$.

This formula shows exactly how Robin closure mixes residual short-end stochastic loading with the boundary-induced linear-exponential component.

10.7 Shape diagnostics of the Robin family

The Robin family has useful short-end diagnostics. Write

$$k = a_R^2, q = k - \rho$$

Then

$$H_R(\tau) = C_0(1 + q\tau)\exp(-k\tau)$$

The boundary value is

$H_R(0) = C_0,$

and the boundary slope is

$H_R'(0) = -C_0\rho.$

Thus, whenever $H_R(0) \neq 0$,

$$\rho = -H_R'(0)/H_R(0)$$

The Robin parameter therefore has a direct observable interpretation: it is the negative boundary slope of the volatility loading relative to its short-end level. Low ρ corresponds to weak anchoring. Larger positive ρ corresponds to stronger suppression of the short-end direction.

The same family also gives a simple turning-point diagnostic. Since

$$H_R'(\tau) = C_0[(q - k) - kq\tau]\exp(-k\tau),$$

an interior critical point satisfies

$$(q - k) - kq\tau = 0$$

If q≠0, then

$$\tau_* = (q - k)/(kq) = -\rho/[k(k - \rho)]$$

The profile has an interior turning point only when this expression is positive and lies in the relevant maturity domain. Thus ρ controls not only the short-end value and boundary slope, but also whether the profile is monotone, shoulder-like, or humped.

This diagnostic is useful for empirical interpretation because it links the estimated boundary behaviour of the volatility loading to the shape of the full maturity profile.

10.8 Sign compatibility and lower-boundary interpretation

The Dirichlet repeated-root loading

$$H_D^{(1)}(\tau) = C_D \tau \exp(-a_D^2 \tau)$$

is non-negative on the half-line when $C_D \geq 0$. It also vanishes at the short end.

For Robin closure, sign compatibility is more delicate. In ρ-form,

$$H_R(\tau) = C_0[1 + (a_R^2 - \rho)\tau] \exp(-a_R^2 \tau)$$

For $C_0{>}0$, non-negativity on the full half-line requires

$$\rho \leq a_R^2$$

If $\rho{>}a_R^2$, the linear factor eventually changes sign. Such a branch may remain mathematically meaningful, but it no longer supports a simple globally non-negative lower-boundary interpretation on the full half-line.

This sign condition is a loading-level admissibility condition. The stronger state-space statement is developed and summarised in Appendix B: under the shifted positive-state KFE structure, the Dirichlet branch admits a shifted inverse-gamma stationary density for each fixed $\tau > 0$, and the lower boundary is inaccessible maturity-by-maturity under the Feller conditions. Robin branches inherit the same local lower-boundary logic when their loading is sign-compatible.

10.9 Comparison with Ritchken–Chuang-type forms

Humped volatility functions of the form

$$\tau \exp(-k\tau)$$

or related linear-exponential variants are familiar in interest-rate modelling. Ritchken–Chuang-type specifications use such maturity profiles because they capture the empirical observation that volatility is often low at the very short end, higher at intermediate maturities, and lower again at long maturities [13, 14].

The IRFR interpretation is different in emphasis.

The resemblance to these specifications does not prove that the r=1 branch is the empirically correct branch. Nor does it show that the IRFR model is uniquely determined by observed humped volatility.

Rather, the point is structural:

a familiar empirical volatility shape appears as an admissible boundary-selected mode.

In conventional modelling, the humped shape is often introduced as a volatility input. In the IRFR framework, the same qualitative shape appears from repeated-root admissibility together with boundary closure.

This provides a structural route to a class of volatility shapes already known to be useful. It does not replace empirical calibration or testing. It motivates them.

The appropriate claim is therefore modest but meaningful: IRFR gives a structural route to familiar humped volatility forms.

10.10 Quadratic models and boundary-selected loadings

Quadratic term-structure models provide an important benchmark for nonlinear interest-rate modelling. Their strategy is to enrich the state representation: yields and bond prices are allowed to depend quadratically on latent factors, and the resulting pricing equations inherit the corresponding polynomial or Riccati-type structure.

The IRFR framework follows a different logic. It does not first introduce quadratic state dependence. Instead, it starts from the rolling-maturity field and asks which maturity loadings are admissible under the interior law and which of those loadings are selected by the boundary condition.

The contrast can be stated simply:

quadratic models generate shape through nonlinear state mapping;

IRFR generates shape through admissible modes and boundary selection.

This distinction matters for the interpretation of humped volatility. In a quadratic or polynomial framework, curvature is produced by the chosen state mapping. In the IRFR framework, the humped component

$$\tau \exp(-a^2\tau)$$

appears as the coalescing-root mode of the maturity-domain admissibility problem and is then selected by the boundary law. The hump is therefore not inserted as a polynomial factor or an exogenous volatility curve. It is the observable survivor of the interior mode space after boundary closure has acted.

This does not make the IRFR framework more general than quadratic term-structure models. It makes a different claim. Before adding polynomial state dependence, one should first account for the geometry of rolling maturity and the selection role of the short-end boundary.

10.11 Representation in HJM coordinates

The boundary-selected IRFR loadings can be translated into conventional forward-rate coordinates through

$$x(t, \tau) = F(t, t + \tau)$$

Equivalently,

$$F(t, T) = x(t, T - t)$$

A maturity loading expressed in τ-coordinates becomes a loading in time-to-maturity coordinates in the HJM representation.

For example, the Dirichlet r=1 loading

$$H(\tau) \propto \tau \exp(-a^2\tau)$$

corresponds to

$$H(T-t) \propto (T-t)exp[-a^2(T-t)]$$

From an HJM perspective, this may appear as an externally specified time-to-maturity volatility function. The HJM drift restriction then gives the corresponding arbitrage-free drift.

From the IRFR perspective, however, the same function has a different origin. It is not first specified as a volatility input. It is selected by the rolling-field boundary law.

This is the representation-versus-explanation distinction introduced earlier. HJM provides the ambient no-arbitrage representation. IRFR explains how rolling-maturity geometry and boundary closure may select particular admissible loadings.

10.12 Empirical interpretation of boundary regimes

The boundary-law framework gives an interpretation of observed volatility profiles in terms of regimes.

A profile with substantial short-end loading is naturally associated with a free or weakly anchored boundary. This corresponds to Neumann/free closure or low-anchoring Robin closure.

A profile with suppressed short-end loading and an interior maximum is naturally associated with a pinned or strongly anchored boundary. The limiting case is Dirichlet closure.

A profile with residual short-end loading and intermediate-maturity curvature is naturally associated with partial anchoring. This corresponds to Robin closure.

Thus empirical volatility shape may contain information about the short-end boundary regime.

This does not mean that observed volatility alone identifies the boundary law uniquely. Different model specifications may produce similar shapes. Multiple factors, stochastic volatility, measurement noise, and market segmentation can all affect observed term-structure volatility.

The claim is more limited:

boundary-selected IRFR modes provide a structural interpretation of volatility shape.

They suggest a way to connect maturity profiles to policy anchoring without treating the volatility function as purely exogenous.

10.13 The role of the r=1 branch

The repeated-root r=1 branch plays a special role because it produces the linear-exponential component

$$\tau \exp(-k\tau)$$

This component is mathematically distinguished: it appears when two admissible exponential modes coalesce. It is also economically interpretable because it gives a natural humped maturity profile.

However, the branch should not be overinterpreted. The empirical usefulness of humped volatility forms does not prove that r=1 is the correct branch. It only shows that the repeated-root branch is a meaningful case to study.

The proper interpretation is conditional:

if the admissible mode structure is near the repeated-root branch,

then boundary selection gives a structural route to humped or partially humped volatility profiles.

Whether observed markets are best described by this branch is an empirical question. It must be tested through calibration, stability, and comparison with alternative branches.

The theoretical contribution is to identify why the r=1 branch is structurally special and how its boundary-selected loadings relate to familiar empirical forms.

10.14 Summary

This chapter has translated the boundary-selected loading structure into observable volatility terms.

The key object is the maturity loading $H(\tau)$. In the one-factor IRFR model,

$$g(\tau) = aH(\tau),$$

and the instantaneous covariance-rate kernel is

$$K(\tau_1, \tau_2) = 2a^2 H(\tau_1)H(\tau_2)$$

Thus $H(\tau)$ determines the maturity shape of volatility.

The Neumann/free benchmark retains the ordinary exponential short-end stochastic loading:

$$H_N(\tau) = A_N \exp(-a_N^2 \tau)$$

The Dirichlet repeated-root branch suppresses short-end loading and selects the humped profile:

$$H_D^{(1)}(\tau) = C_D \tau \exp(-a_D^2 \tau)$$

The Robin repeated-root branch gives the partially anchored mixed profile:

$$H_R(\tau) = (C_0 + C_1 \tau) \exp(-a_R^2 \tau)$$

These are not arbitrary volatility functions. They are boundary-selected maturity loadings. Their empirical relevance must be tested, but their structural origin is clear.

observable volatility shape is the footprint of boundary selection.

Chapter 11

Stochastic Long-Maturity Boundary States

The preceding chapters classify admissible IRFR fields through their short-end boundary behaviour. Neumann/free, Dirichlet, and Robin/mixed boundary regimes determine how the maturity field behaves at the instantaneous short end. This chapter completes the boundary-field picture by relaxing the opposite end of the maturity manifold. Instead of forcing the long-maturity boundary to coincide pathwise with the fixed equilibrium anchor x_∞, we distinguish the equilibrium anchor from the realised asymptotic boundary state.

The central distinction is

$$x_\infty \quad \text{as the equilibrium long-run anchor,}$$

and

$$X_t \quad \text{as the realised long-maturity boundary state.}$$

The boundary displacement is

$$Y_t := X_t - x_\infty.$$

The construction does not introduce a second Brownian risk factor. The displacement Y_t is carried by the same admissible IRFR risk direction as the finite-maturity field. In this sense, stochastic long-boundary motion is a boundary excitation of the existing one-factor field, not an additional independent source of randomness.

11.1 Motivation

In the fixed long-boundary formulation, the maturity field is anchored asymptotically at x_∞. Finite-maturity deviations may be transported along the curve, but the limiting long end remains fixed:

$$\lim_{\tau\to\infty} x_{t,T} = x_\infty, \qquad \tau = T - t.$$

This is appropriate when the long-run equilibrium anchor is treated as both the economic regime reference and the realised long-end state. It is too restrictive, however, if persistent long-run level shifts or long-end volatility are to be represented within the same boundary-field architecture.

The stochastic long-boundary extension separates these roles. The parameter x_∞ remains the equilibrium anchor. The process X_t is the realised long-end boundary state. When $X_t \neq x_\infty$, the asymptotic boundary of the field is displaced from equilibrium. This displacement is not modelled by adding a second factor. Instead, it is loaded onto the same admissible risk direction that already governs the IRFR field.

The purpose of this chapter is therefore to show that stochastic long-end boundary motion can be incorporated without abandoning the single-risk-direction structure. The short-end boundary classification remains intact. The long-boundary layer is orthogonal to it.

11.2 Boundary decomposition

Let

$$\tau = T - t.$$

For a given short-end boundary branch j, define the branch transport reference

$$b_j(\tau) = x_\infty \left(1 - e^{-r_j a^2 \tau}\right).$$

Here r_j is the branch parameter associated with the short-end boundary regime. The stochastic long-boundary loading is defined by

$$L(\tau) = 1 - e^{-a^2 \tau}.$$

This loading satisfies

$$L(0) = 0, \qquad \lim_{\tau \to \infty} L(\tau) = 1.$$

The stochastic long-boundary decomposition is

$$x^j_{t,T} = b_j(\tau) + L(\tau)(X_t - x_\infty) + Z^j_{t,T}.$$

Equivalently,

$$Z^j_{t,T} := x^j_{t,T} - b_j(\tau) - L(\tau)(X_t - x_\infty).$$

The residual $Z^j_{t,T}$ is the finite-maturity curve-shape component after removing both the branch transport reference and the transported long-boundary displacement.

At the short end,

$$b_j(0) = 0, \qquad L(0) = 0,$$

so

$$x^j_{t,t} = Z^j_{t,t}.$$

At the long end, provided $Z^j_{t,T} \to 0$ as $\tau \to \infty$,

$$\lim_{\tau \to \infty} x^j_{t,T} = x_\infty + (X_t - x_\infty) = X_t.$$

Thus the decomposition preserves the short-end state while replacing the fixed long-end limit x_∞ with the realised stochastic boundary state X_t.

11.3 Single-risk-direction admissibility

The stochastic boundary branch is defined by requiring the full displacement from the branch transport reference to be carried by the unique admissible IRFR risk direction:

$$x^j_{t,T} - b_j(\tau) = L(\tau)(X_t - x_\infty) + Z^j_{t,T}.$$

The boundary displacement and the finite-maturity residual therefore do not generate separate Brownian directions. They enter the same drift–diffusion loading.

Define

$$f_j(t,T) = a^2 \left[x^j_{t,T} - b_j(\tau)\right],$$

and

$$g_j(t,T) = a \left[x^j_{t,T} - b_j(\tau)\right].$$

Then, whenever the effective loading is non-zero,

$$\frac{f_j(t,T)}{g_j(t,T)} = a.$$

The physical market signal-to-noise ratio is therefore preserved:

$$\mathrm{MSNR}^{\mathbb{P}}_j = \frac{f_j(t,T)}{\sqrt{2}\, g_j(t,T)} = \frac{a}{\sqrt{2}}.$$

Under the IRFR field convention,

$$dx^j_{t,T} = -f_j(t,T)\, dt + \sqrt{2}\, g_j(t,T)\, dW_t,$$

the stochastic long-boundary branch becomes

$$dx^j_{t,T} = -a^2 \left[x^j_{t,T} - b_j(\tau)\right] dt + \sqrt{2}a \left[x^j_{t,T} - b_j(\tau)\right] dW_t.$$

The stochastic boundary changes the level of the effective loading, but it does not change the admissible risk direction.

11.4 Asymptotic boundary dynamics

Taking the long-maturity limit of the field equation gives the boundary-mode dynamics. Since

$$x^j_{t,T} \to X_t, \qquad b_j(\tau) \to x_\infty \quad \text{as } \tau \to \infty,$$

the limiting equation is

$$dX_t = -a^2(X_t - x_\infty)\, dt + \sqrt{2}a(X_t - x_\infty)\, dW_t.$$

Equivalently, for

$$Y_t := X_t - x_\infty,$$

we have

$$dY_t = -a^2 Y_t\, dt + \sqrt{2}a Y_t\, dW_t.$$

A boundary mode, in the sense used here, is the limiting maturity-field excitation carried by the asymptotic boundary state rather than by an interior curve-shape residual. The long-boundary displacement Y_t is therefore a multiplicative boundary mode. For $Y_{t_0} \neq 0$,

$$Y_t = Y_{t_0} \exp\left(-2a^2(t - t_0) + \sqrt{2}a(W_t - W_{t_0})\right).$$

The sign of the displacement is preserved. If $Y_{t_0} = 0$, then $Y_t = 0$ thereafter and $X_t = x_\infty$. Thus, in this multiplicative long-boundary subclass, departures from the equilibrium long-end anchor must either be present as an initial boundary displacement or introduced through a richer additive or regime-switching boundary mechanism.

This is a modelling implication of the multiplicative construction. The stochastic long boundary represents persistence and propagation of an existing long-run displacement. It does not by itself generate a new displacement from the exactly anchored state $X_t = x_\infty$.

11.5 Generator corroboration

The same boundary law is recovered from the limiting infinitesimal generator. In the long-maturity limit, write

$$y = X - x_\infty.$$

For smooth test functions φ, the physical generator associated with the boundary process is

$$(\mathcal{A}\varphi)(X) = -a^2(X - x_\infty)\varphi_X(X) + a^2(X - x_\infty)^2\varphi_{XX}(X).$$

Applying $\mathcal{A}$ to $\varphi(X) = X$ gives the drift

$$-a^2(X - x_\infty).$$

The second-order term identifies the diffusion variance. For $X_0 = X$,

$$\mathcal{A}\left[(X - X_0)^2\right]\Big|_{X=X_0} = 2a^2(X_0 - x_\infty)^2.$$

Thus

$$\sigma^2(X) = 2a^2(X - x_\infty)^2,$$

with signed one-risk-direction choice

$$\sigma(X) = \sqrt{2}a(X - x_\infty).$$

The generator calculation therefore corroborates the asymptotic boundary SDE without introducing a second Brownian source or changing the stochastic basis.

11.6 Expected boundary reversion

For fixed maturity T, let

$$m_j(t, T) := \mathbb{E}[x^j_{t,T}].$$

The conditional mean satisfies

$$\frac{dm_j(t,T)}{dt} = -a^2\left[m_j(t,T) - b_j(T - t)\right].$$

The equilibrium mean curve in maturity coordinates is

$$x_e^j(\tau) = x_\infty - \frac{x_\infty}{1+r_j} e^{-r_j a^2 \tau}.$$

Thus,

$$\mathbb{E}[x_{t,T}^j] = x_e^j(T-t) + e^{-a^2(t-t_0)} \left[x_{t_0,T}^j - x_e^j(T-t_0) \right].$$

Taking the long-maturity limit gives

$$\lim_{T-t\to\infty} \mathbb{E}[x_{t,T}^j] = x_\infty + (X_{t_0} - x_\infty) e^{-a^2(t-t_0)}.$$

The expected realised long boundary therefore mean-reverts toward the equilibrium anchor x_∞ at the physical adjustment speed a^2.

The branch transport reference $b_j(\tau)$ and the equilibrium mean curve $x_e^j(\tau)$ should not be confused. The former is the instantaneous reference appearing inside the local drift loading. The latter is the stationary first-moment profile after accounting for maturity ageing, $\tau = T - t$. Except in the degenerate case $r_j = 0$, these two objects differ.

11.7 Long-end variance floor

The instantaneous variance rate of the stochastic long-boundary field is

$$v_j(t,T) := \frac{d}{dt} \operatorname{Var}_t(x_{t,T}^j) = 2a^2 \left[x_{t,T}^j - b_j(\tau) \right]^2.$$

Taking the long-maturity limit yields

$$\lim_{\tau\to\infty} v_j(t,T) = 2a^2 (X_t - x_\infty)^2.$$

This is the main observable implication of the stochastic long-boundary extension. In the fixed-boundary model, the long-end loading vanishes and the local variance collapses to zero. In the stochastic-boundary model, long-end variance remains positive whenever the realised long boundary differs from the equilibrium anchor.

The long-end variance floor is therefore tied to three interpretable quantities:

$$a^2, \qquad x_\infty, \qquad X_t.$$

The adjustment speed a^2 controls propagation and reversion. The anchor x_∞ represents the equilibrium long-run regime. The realised boundary X_t represents the market-implied long-run state. Persistent long-end volatility is therefore represented by the distance between the realised long-run boundary and the equilibrium regime anchor, rather than by an additional curve-fitting factor.

11.8 Curve representation

Impose the multiplicative residual maturity-loading Ansatz

$$Z_{t,T}^j = x_{t,t}^j e^{-a^2 \tau}.$$

Substituting into the stochastic boundary decomposition gives

$$x_{t,T}^{j} = x_\infty \left(1 - e^{-r_j a^2 \tau}\right) + \left(1 - e^{-a^2\tau}\right)(X_t - x_\infty) + x_{t,t}^{j} e^{-a^2\tau}.$$

Equivalently,

$$x_{t,T}^{j} = x_{t,t}^{j} e^{-a^2\tau} + X_t \left(1 - e^{-a^2\tau}\right) + x_\infty \left(e^{-a^2\tau} - e^{-r_j a^2 \tau}\right).$$

The case $r_j = 1$ is the unique bridge case in which the explicit x_∞ loading cancels from the static cross-sectional representation. The curve becomes

$$x_{t,T}^{j} = X_t + \left(x_{t,t}^{j} - X_t\right) e^{-a^2\tau}.$$

In this case, the observed curve is an exponential bridge between the realised short-end state and the realised long-end boundary state. The equilibrium anchor x_∞ still enters through the boundary dynamics, expected reversion, and long-end variance floor, but it no longer appears as a separate visible loading in the cross-sectional curve.

11.9 Branch specialisations

The stochastic long-boundary machinery is common to all short-end boundary branches. The branch label j changes the short-end admissibility condition and the parameter r_j appearing in the transport reference $b_j(\tau)$. It does not change the one-risk-direction boundary law, the MSNR preservation result, or the long-end variance floor.

For the Neumann/free branch,

$$r_j = r.$$

For the Dirichlet branch,

$$r_D = \frac{\lambda_D^2}{a^2}.$$

For the Robin/mixed branch,

$$r_R = \frac{\lambda_R^2}{a^2}.$$

Across all three branches,

$$b_j(\tau) = x_\infty \left(1 - e^{-r_j a^2 \tau}\right),$$

and the stochastic long-boundary curve representation remains

$$x_{t,T}^{j} = x_\infty \left(1 - e^{-r_j a^2 \tau}\right) + \left(1 - e^{-a^2\tau}\right)(X_t - x_\infty) + x_{t,t}^{j} e^{-a^2\tau}.$$

The Dirichlet and Robin amplitudes are not inherited from the Neumann/free branch. They are fixed by the branch-specific equilibrium or moment-transport calculation developed in the short-end boundary chapters. In the present chapter, the long-boundary mechanism is layered on top of those branch-specific short-end closures.

11.10 Relation to risk-neutral pricing

The stochastic long-boundary extension is a physical-field construction. It identifies how a realised long-maturity boundary state can move within the same admissible IRFR risk direction. The following chapter introduces the risk-neutral pricing layer.

This ordering is important. The money-market account is generated by the realised short-end roll $x_{t,t}$. The stochastic long boundary X_t does not directly fund the bank account. Instead, it enters bond prices through the IRFR strip

$$\{x_{t,u} : u \in [t, T]\}.$$

When the long boundary is stochastic, the bond-price functional inherits dependence on the realised asymptotic state. A full no-arbitrage treatment must therefore trace X_t through the risk-neutral drift restriction and the rolling-maturity transport equation.

The present chapter does not impose that pricing restriction. It completes the physical boundary-field architecture. The next chapter adds the money-market numeraire, the martingale condition for discounted bond prices, and the change from the physical measure to the pricing measure.

11.11 Summary

The fixed-boundary IRFR model identifies the long end with the equilibrium anchor x_∞. The stochastic long-boundary extension separates this anchor from the realised long-end state X_t. The resulting displacement

$$Y_t = X_t - x_\infty$$

is carried by the same admissible IRFR risk direction as the finite-maturity field.

The construction preserves the physical loading relation

$$\frac{f_j}{g_j} = a,$$

and hence preserves the one-risk-direction MSNR structure. The asymptotic boundary dynamics are

$$dX_t = -a^2(X_t - x_\infty)\,dt + \sqrt{2}a(X_t - x_\infty)\,dW_t.$$

The long-end variance floor is

$$\lim_{\tau\to\infty} v_j(t,T) = 2a^2(X_t - x_\infty)^2.$$

Thus the monograph's boundary-field architecture has two complementary layers. Short-end boundary regimes determine admissibility at $\tau = 0$. The stochastic long-boundary state determines persistent asymptotic displacement as $\tau \to \infty$. Together they allow the IRFR field to represent both local short-rate boundary behaviour and long-run regime displacement without adding an independent Brownian factor.

Chapter 12

Bond Pricing, Rolling Replication, and the Risk-Neutral IRFR Measure

12.1 Introduction

The preceding chapters develop IRFR as a theory of admissible maturity fields, boundary-selected modes, and observable volatility profiles. By the end of Chapter 10, the framework has a field equation, a boundary classification, and a set of maturity loadings selected by Neumann, Dirichlet, or Robin closure. What remains is the pricing bridge from admissible local rate-field coordinates to traded bond prices.

This chapter supplies that bridge. The key point is that the money-market account is generated only by the realised short-end roll. The rest of the IRFR strip is not itself a bank account. Non-short IRFRs are locally zero-value forward-rate contracts: they do not require initial funding and do not directly accrue the cash account at inception. Their role is to represent, hedge, or replicate future rolling short-rate exposures.

Risk-neutral pricing is therefore imposed in the standard way, but on the right object. The money-market account is the numeraire generated by the realised short rate. Zero-coupon bond prices are generated by integrating the current IRFR strip. The no-arbitrage condition is that these bond prices, when discounted by the short-rate money-market account, are martingales under the pricing measure.

The resulting construction is an IRFR-constrained HJM framework: IRFR restricts the admissible maturity geometry of the field and its volatility modes, while the risk-neutral measure supplies the drift required for arbitrage-free bond pricing [6, 12]. The chapter does not claim to complete derivative pricing for bond options, caps, floors, or swaptions. It establishes the bond-pricing layer and identifies the remaining derivative-pricing problem as the law under $\mathbb{Q}$ of integrated IRFR field functionals.

12.2 The money-market account at an initial date

Begin at a fixed date t_0, and consider a small time step Δt. The realised short-end field value is

$$r_{t_0} = x_{t_0,t_0}.$$

This is the rate that accrues the cash account over the first interval $[t_0, t_0 + \Delta t]$. In simple-discrete notation,

$$B_{t_0+\Delta t} = B_{t_0}\left(1 + x_{t_0,t_0}\Delta t\right),$$

and in continuous time the corresponding differential form is

$$dB_t = x_{t,t}B_t\, dt.$$

The non-short IRFR

$$x_{t_0,t_0+\Delta t}$$

has a different status. It is not a cash-account accrual rate at time t_0. It is a locally zero-value forward-rate exposure to the next short-rate roll. It can be used to represent the second-period rolling short-rate exposure, but it does not itself fund or accrue the bank account at inception.

Over two short periods, the realised rolling cash-account return is built from the sequence

$$x_{t_0,t_0} \quad \text{and} \quad x_{t_0+\Delta t,t_0+\Delta t}.$$

At time t_0, the second object has not yet been realised. The current IRFR

$$x_{t_0,t_0+\Delta t}$$

is the locally zero-value contract representing that next roll. Thus, to first order in Δt, the two-period rolling exposure can be written at t_0 as

$$x_{t_0,t_0}\Delta t + x_{t_0,t_0+\Delta t}\Delta t.$$

This expression should not be read as saying that both terms accrue the cash account at t_0. Only x_{t_0,t_0} does. The second term is the current IRFR representation of the next short-rate roll.

This elementary two-period construction is the pricing intuition of the chapter. The cash account is generated by the realised short end. The remaining IRFR strip is a strip of locally zero-value contracts used to represent future short-rate rolls.

12.3 Generalisation to an arbitrary time

At a general time t, define the realised short rate by

$$r_t := x_{t,t}.$$

The money-market account is therefore

$$dB_t = r_t B_t\, dt = x_{t,t}B_t\, dt, \qquad B_0 = 1,$$

so that

$$B_t = \exp\left(\int_0^t x_{s,s}\, ds\right).$$

This definition deliberately assigns funding status only to the realised short-end roll. The forward IRFRs

$$x_{t,T}, \qquad T > t,$$

or, in rolling-maturity notation,

$$x(t, \tau), \qquad \tau > 0,$$

are not components of the money-market account at time t. They are locally zero-value contracts representing the current strip of future rolling short-rate exposures.

The distinction is essential. The money-market account is a self-financing numeraire. A non-short IRFR is a rate-field coordinate or forward-rate contract. It may determine the price of a traded bond once integrated across maturity, but it is not itself a strictly positive tradable price process and cannot itself serve as a numeraire.

12.4 The IRFR strip and zero-coupon bond prices

Let $f(t, T)$ denote the instantaneous forward rate observed at calendar time t for calendar maturity $T \geq t$. The IRFR field is the same forward-rate surface written in rolling-maturity coordinates:

$$\tau = T - t, \qquad x(t, \tau) = f(t, t + \tau).$$

Equivalently,

$$f(t, T) = x(t, T - t).$$

A zero-coupon bond price is obtained from the current forward-rate strip:

$$P(t, T) = \exp\left(-\int_t^T f(t, u)\, du\right).$$

In IRFR coordinates this becomes

$$P(t, T) = \exp\left(-\int_0^{T-t} x(t, s)\, ds\right).$$

This formula uses the current IRFR strip to represent the price of future rolling short-rate exposure. It does not say that the whole strip accrues the cash account. Rather, the strip prices the traded bond whose payoff is one unit at T, while the money-market account is generated only by the realised short-end path $x_{s,s}$.

The relationship between the two objects is the standard no-arbitrage bridge:

$$P(t, T) = \mathbb{E}_t^{\mathbb{Q}}\left[\exp\left(-\int_t^T x_{s,s}\, ds\right)\right],$$

where $\mathbb{Q}$ is the risk-neutral measure associated with the money-market account. The left-hand side is the current strip valuation. The right-hand side is the discounted expectation of future realised short-rate rolls.

12.5 Stochastic consistency of rolling replication

The generalisation from t_0 to arbitrary t introduces the central stochastic consistency issue. The future realised short rate

$$x_{t+\Delta t, t+\Delta t}$$

is not pathwise known at time t, and it is not equal in advance to the current forward object

$$x_{t,t+\Delta t}.$$

The latter is the time-t locally zero-value representation of the next short-rate roll. The former is the subsequently realised short-end state of the field.

They are linked by the stochastic IRFR dynamics. The current IRFR strip is therefore a present representation of future rolling cash-account exposures, but the future rolls themselves are stochastic. This is why risk-neutral pricing cannot be imposed by treating every IRFR as a cash-account rate. It must instead be imposed as a martingale condition on bond prices generated by the IRFR strip.

Thus the correct no-arbitrage condition is

$$\frac{P(t,T)}{B_t} \quad \text{is a } \mathbb{Q}\text{-martingale.}$$

Equivalently, for $s \geq t$,

$$\mathbb{E}_t^{\mathbb{Q}}\left[\frac{P(s,T)}{B_s}\right] = \frac{P(t,T)}{B_t}.$$

This condition equates the current strip-generated bond price with the pricing-measure value of the future realised rolling short-rate cash flows.

12.6 Physical and pricing dynamics

Under the physical measure $\mathbb{P}$, write the local IRFR field dynamics as

$$dx_{t,T} = \mu^{\mathbb{P}}(t,T)\,dt + \sigma(t,T)\,dW_t^{\mathbb{P}}.$$

The monograph uses the IRFR sign convention introduced in the field-dynamics chapters: the deterministic pull of the field is written as $-f\,dt$, and the diffusion coefficient is written as $\sqrt{2}\,g$. Thus

$$\mu^{\mathbb{P}}(t,T) = -f(t,T), \qquad \sigma(t,T) = \sqrt{2}\,g(t,T).$$

This is a convention for the IRFR field variable, not a change in the standard HJM money-market numeraire construction.

Let the change of measure be given by

$$dW_t^{\mathbb{Q}} = dW_t^{\mathbb{P}} + \lambda_t\,dt,$$

where λ_t is the market price of IRFR field risk. Then

$$dW_t^{\mathbb{P}} = dW_t^{\mathbb{Q}} - \lambda_t\,dt,$$

and hence

$$dx_{t,T} = \left[\mu^{\mathbb{P}}(t,T) - \lambda_t\sigma(t,T)\right]dt + \sigma(t,T)dW_t^{\mathbb{Q}}.$$

Substituting the IRFR notation gives

$$dx_{t,T} = -\left[f(t,T) + \sqrt{2}\lambda_t g(t,T)\right]dt + \sqrt{2}g(t,T)dW_t^{\mathbb{Q}}.$$

Therefore the risk-neutral drift loading is

$$f^{\mathbb{Q}}(t,T) = f(t,T) + \sqrt{2}\lambda_t g(t,T).$$

The physical Market Signal-to-Noise Ratio is the signed drift loading divided by the diffusion loading in the established IRFR convention:

$$\mathrm{MSNR}^{\mathbb{P}}(t,T) := \frac{f(t,T)}{\sqrt{2}\, g(t,T)}.$$

Under the physical IRFR admissibility relation

$$\frac{f(t,T)}{g(t,T)} = a,$$

this becomes

$$\mathrm{MSNR}^{\mathbb{P}} = \frac{a}{\sqrt{2}}.$$

For the Chapter 5 branch parameterisation $r = \lambda^2/a^2$, with λ denoting the branch spectral parameter, this may also be written as

$$\mathrm{MSNR}^{\mathbb{P}} = \frac{\lambda}{\sqrt{2r}}.$$

Under the pricing measure,

$$\frac{f^{\mathbb{Q}}(t,T)}{g(t,T)} = a + \sqrt{2}\lambda_t.$$

Thus the one-dimensional Brownian risk direction is preserved by the change of measure. The physical MSNR is transformed by the market price of risk through the drift loading assigned to the same volatility direction.

For the simplest first-pass pricing layer one may take $\lambda_t = \lambda$ constant. Then

$$a_{\mathbb{Q}} = a + \sqrt{2}\lambda,$$

and the risk-neutral dynamics remain in the same one-risk-direction family with a shifted drift-to-volatility ratio. State-dependent or time-dependent market prices of risk can be introduced later, but they belong to calibration and empirical implementation rather than to the basic money-market-account construction.

12.7 Relation to the HJM drift condition

The preceding measure-change calculation describes how the physical IRFR drift is adjusted by the market price of risk. The martingale condition on $P(t,T)/B_t$ supplies the additional no-arbitrage restriction on the pricing-measure drift of the forward-rate curve.

In conventional instantaneous-forward-rate notation, with calendar maturity T, write

$$df(t,T) = \alpha^{\mathbb{Q}}(t,T)\, dt + \sigma_f(t,T) dW_t^{\mathbb{Q}}.$$

The one-factor HJM drift restriction is

$$\alpha^{\mathbb{Q}}(t,T) = \sigma_f(t,T) \int_t^T \sigma_f(t,u)\, du,$$

under the bond-price convention

$$P(t,T) = \exp\left(-\int_t^T f(t,u)\,du\right).$$

The IRFR field is normally written in rolling-maturity coordinates,

$$\tau = T - t, \qquad x(t,\tau) = f(t, t+\tau).$$

Differentiation at fixed calendar maturity and differentiation at fixed rolling maturity differ by the transport term:

$$\partial_t f(t,T)|_T = \partial_t x(t,\tau)|_\tau - \partial_\tau x(t,\tau), \qquad \tau = T - t.$$

Therefore, if the rolling-coordinate $\mathbb{Q}$-dynamics are written as

$$dx(t,\tau) = \widetilde{\mu}^{\mathbb{Q}}(t,\tau)\,dt + \sigma(t,\tau)dW_t^{\mathbb{Q}},$$

then the calendar-maturity drift is

$$\alpha^{\mathbb{Q}}(t,t+\tau) = \widetilde{\mu}^{\mathbb{Q}}(t,\tau) - \partial_\tau x(t,\tau).$$

The HJM restriction in rolling-maturity coordinates is consequently

$$\widetilde{\mu}^{\mathbb{Q}}(t,\tau) - \partial_\tau x(t,\tau) = \sigma(t,\tau)\int_0^\tau \sigma(t,s)\,ds,$$

or equivalently

$$\boxed{\widetilde{\mu}^{\mathbb{Q}}(t,\tau) = \partial_\tau x(t,\tau) + \sigma(t,\tau)\int_0^\tau \sigma(t,s)\,ds.}$$

Substituting the IRFR volatility convention $\sigma(t,\tau) = \sqrt{2}\,g(t,\tau)$ gives

$$\widetilde{\mu}^{\mathbb{Q}}(t,\tau) = \partial_\tau x(t,\tau) + 2g(t,\tau)\int_0^\tau g(t,s)\,ds.$$

In the remaining displays of this subsection, $f(t,\tau)$ denotes the IRFR drift-loading coefficient from the field-dynamics chapters, not the forward-rate field $f(t,T)$ used above in calendar-maturity notation. The IRFR $\mathbb{Q}$-dynamics write the same rolling-coordinate drift as

$$dx(t,\tau) = -f^{\mathbb{Q}}(t,\tau)\,dt + \sqrt{2}g(t,\tau)dW_t^{\mathbb{Q}}.$$

Hence no-arbitrage imposes

$$\boxed{f^{\mathbb{Q}}(t,\tau) = -\partial_\tau x(t,\tau) - 2g(t,\tau)\int_0^\tau g(t,s)\,ds.}$$

Combining this with the measure-change identity from the previous section therefore yields the pricing consistency condition

$$\boxed{f(t,\tau) + \sqrt{2}\lambda_t g(t,\tau) = -\partial_\tau x(t,\tau) - 2g(t,\tau)\int_0^\tau g(t,s)\,ds.}$$

This is the transport-corrected HJM restriction in IRFR coordinates. IRFR admissibility supplies the maturity geometry and volatility loading; the martingale pricing condition supplies the $\mathbb{Q}$-drift compatible with bond-market no-arbitrage.

This separation is important. IRFR admissibility alone does not prove self-financing no-arbitrage. It constrains the field's admissible maturity geometry. The martingale pricing condition is the additional pricing restriction that turns the field into a bond-pricing model.

12.8 Boundary interpretation of the short rate

The short rate entering the money-market account is the realised short-end boundary value

$$r_t = x_{t,t}.$$

This is consistent with the boundary-field interpretation of the earlier chapters. The short end is not an afterthought appended for pricing. It is already the boundary of the rolling-maturity field.

The pricing layer therefore does not erase the boundary interpretation. It uses it. The money-market account is generated by the realised boundary roll; non-short IRFRs represent future boundary rolls through the current forward strip; and bond prices are the traded instruments produced by integrating that strip across maturity.

This also clarifies the status of boundary laws. Neumann/free, Dirichlet, and Robin regimes restrict the admissible maturity loading of the field. The risk-neutral measure does not replace those boundary conditions. It changes the drift needed for traded bond prices to be martingales after discounting by the short-rate numeraire.

A pricing-relevant consequence arises under Dirichlet closure. If the short-end boundary is fixed at x_D, then the money-market account is conditionally deterministic:

$$\boxed{\text{Dirichlet closure:} \qquad r_t = x_D, \qquad B_t = e^{x_D t} \quad \text{conditionally deterministic.}}$$

The remaining curve may still contain non-short IRFR exposure, but the funding numeraire itself is pinned by the fixed short-end boundary.

In the stochastic long-boundary extension, the realised long-end state also affects the field and hence the bond-price functional. But it does not directly fund the money-market account. The cash account remains a short-end object generated by $x_{t,t}$. A full no-arbitrage treatment of the stochastic long-boundary case must trace the moving asymptotic state through the pricing PDE and through the integrated forward-rate strip; that extension is deferred to the long-boundary pricing analysis.

12.9 Derivative-pricing extension

Once the zero-coupon bond price is written as

$$P(t, t+\tau) = \exp\left(-\int_0^\tau x(t,s)\,ds\right),$$

any bond derivative becomes a contingent claim on an integrated IRFR field. For example, if a reduced one-factor representation has the form

$$x(t,\tau) = x_0(\tau) + Z_t H(\tau),$$

then

$$P(t, t+\tau) = \exp\left(-\int_0^\tau x_0(s)\,ds - Z_t \int_0^\tau H(s)\,ds\right).$$

The admissible maturity loading H determines bond-price sensitivity through the integrated loading

$$\int_0^\tau H(s)\,ds.$$

As derived in Chapter 5 for the repeated-root humped branch,

$$H(\tau) = C\tau e^{-\alpha\tau},$$

the integrated loading is

$$\int_0^\tau H(s)\,ds = \frac{C}{\alpha^2}\left[1 - e^{-\alpha\tau}(1+\alpha\tau)\right].$$

Hence the field-to-bond mapping is explicit. The remaining derivative-pricing problem is to specify the risk-neutral law of the relevant factor or integrated field functional under $\mathbb{Q}$, then price the payoff under the chosen numeraire.

This chapter therefore establishes bond-pricing consistency but does not claim to complete the pricing of options on bonds, caps, floors, or swaptions. Those require additional assumptions about factor reduction, market completeness, payoff-specific numeraires, numerical methods, and empirical calibration of λ_t.

12.10 What this construction does and does not prove

This chapter establishes a bond-pricing interpretation for IRFR. It shows that:

1. the money-market account is generated only by the realised short-end roll $x_{t,t}$;
2. non-short IRFRs are locally zero-value forward-rate contracts rather than bank-account accrual rates;
3. the current IRFR strip represents future rolling short-rate exposures;
4. zero-coupon bond prices are generated by integrating the IRFR strip;
5. risk-neutral pricing is imposed by requiring $P(t,T)/B_t$ to be a $\mathbb{Q}$-martingale;
6. the change from $\mathbb{P}$ to $\mathbb{Q}$ preserves the single Brownian risk direction but changes the drift loading through the market price of risk;
7. the physical MSNR is transformed by the market price of risk rather than destroyed by the change of measure;
8. IRFR admissibility and no-arbitrage martingale pricing are complementary restrictions, not identical ones.

The construction does not prove uniqueness of the risk-neutral measure, market completeness, or closed-form derivative prices. Nor does it imply that IRFR admissibility alone enforces no-arbitrage. The field equation and boundary regime restrict admissible maturity geometry. The martingale condition supplies the pricing drift.

The main conclusion is therefore precise. The money-market account is not the integral of the whole IRFR curve. It is the accumulated realised short-end roll. The rest of the IRFR curve is a strip of locally zero-value forward-rate contracts representing future short-rate rolls. Risk-neutral pricing is the consistency condition equating the current strip-generated bond price with the $\mathbb{Q}$-discounted expectation of future realised short-rate accruals.

Chapter 13

Calibration, Identification, and Empirical Discipline

13.1 Introduction

The preceding chapters derived the IRFR framework as a boundary-value theory of interest-rate dynamics.

The early chapters introduced the instantaneous rolling forward rate as a maturity-domain field. The middle chapters derived the affine drift operator, admissible maturity modes, and the r-branch. The boundary chapters showed how Neumann/free, Dirichlet, and Robin closures select different realised maturity loadings. Chapter 10 then translated those boundary-selected loadings into observable volatility profiles.

The present chapter turns from structural derivation to empirical discipline.

The purpose is not to claim that the IRFR framework has already been empirically validated. It has not. The purpose is to state how the framework should be taken to data, what must be estimated, what can and cannot be inferred from volatility shape, and how boundary regimes might be tested.

Estimation and identification are distinct tasks. Estimation concerns recovering the empirical maturity loading $H(\tau)$, and possibly its scale, from data. Identification concerns mapping features of that loading — its decay, short-end value, short-end slope, hump location, and factor structure — to structural parameters such as k, ρ, and the boundary regime.

There is also a preliminary data-construction step. The IRFR field is not observed directly. Market data are typically available as bond prices, swap rates, futures, deposits, OIS curves, or option-implied quantities. Recovering $x(t, \tau)$ and estimating changes in $H(\tau)$ therefore requires curve construction, interpolation, smoothing, and a choice of instruments and discounting conventions. These choices are not innocuous. They affect the estimated maturity loading and, consequently, any inferred boundary parameter. Empirical IRFR calibration should therefore treat curve construction as part of the modelling procedure rather than as a neutral preprocessing step.

The guiding principle is:

empirical fit is evidence, not proof of canonicality.

A humped volatility profile may be consistent with the repeated-root r = 1 branch, but it does not prove that the repeated-root branch is the true branch. A good calibration may support the usefulness of the framework, but it does not establish uniqueness. A boundary-selected loading may explain a volatility shape, but the same shape may be approximated by other models.

This chapter therefore sets out ten principles for calibration, identification, and empirical testing. These principles are not additional axioms. They are empirical safeguards. They are intended to prevent the structural claims of the monograph from being overstated when the model is confronted with data.

The ten principles are:

1. Estimate the maturity loading, not only the volatility level.
2. Separate stochastic dimension from maturity geometry.
3. Identify the decay scale.
4. Estimate boundary anchoring from short-end behaviour.
5. Do not infer $r = 1$ from hump resemblance alone.
6. Distinguish scale from shape.
7. Use principal components carefully.
8. Treat boundary regimes as estimable hypotheses.
9. Check positivity and admissibility separately.
10. Treat empirical fit as evidence, not proof.

Together, these principles define the empirical programme of the IRFR framework.

13.2 Principle 1: estimate the loading, not only the volatility level

The primary empirical object in the one-factor IRFR model is the maturity loading

$$H(\tau)$$

Under the convention used in the previous chapters,

$$g(\tau) = aH(\tau),$$

and the conventional volatility coefficient is

$$\sqrt{2}aH(\tau)$$

The overall volatility level matters, but it is not the main structural object. The structural object is the shape of $H(\tau)$: where the curve is sensitive to the common shock, how that sensitivity changes across maturity, and how the loading behaves at the short end.

A calibration that matches only the total variance level is therefore insufficient. It may reproduce the average magnitude of rate changes while missing the maturity-domain geometry.

The IRFR framework makes a stronger claim. It says that the shape of the maturity loading contains information about the boundary law. A loading with substantial short-end value suggests a different boundary regime from a loading that vanishes at the short end. A loading with a humped component suggests a different admissible-mode structure from a purely monotone exponential decay.

Thus calibration must estimate not only how much volatility there is, but where that volatility lives along the maturity domain.

The empirical question is therefore:

what is the estimated shape of $H(\tau)$?

Only after this shape is estimated can one ask which IRFR boundary regime it resembles.

13.3 Principle 2: separate stochastic dimension from maturity geometry

A one-factor stochastic model does not imply a flat or trivial maturity profile.

In the one-factor IRFR model, all maturities share a common Brownian shock,

dW_t.

The covariance kernel has the rank-one form

$$K(\tau_1, \tau_2) = 2a^2 H(\tau_1) H(\tau_2)$$

The stochastic dimension is one. But the deterministic maturity geometry is encoded in $H(\tau)$, and this function may be non-trivial.

This distinction is essential. A model may be one-factor in randomness while still having structured maturity sensitivity. The curve may react to one shock, but the loading of that shock across maturities can be humped, shouldered, monotone, or boundary-active.

Empirically, this means that the number of stochastic factors and the shape of the volatility loading should not be conflated.

A one-factor PCA approximation, for example, may identify a dominant empirical volatility direction. That direction is not merely a scalar factor. It is a function over maturity. In the one-factor IRFR setting, that function is the empirical counterpart of $H(\tau)$.

Under the one-factor IRFR restriction, the leading empirical eigenfunction must coincide, up to scale and estimation error, with the maturity loading $H(\tau)$. This is a testable restriction, not merely an interpretation.

Thus the empirical task is not only to ask how many factors are needed. It is also to ask whether the leading maturity loading has a shape compatible with an IRFR boundary-selected mode.

13.4 Principle 3: identify the decay scale

The admissible IRFR loadings contain a maturity-decay scale.

In the repeated-root branch, the generic loading has the form

$$H(\tau) = (C_1 + C_2\tau)e^{-k\tau}$$

In the $r = 1$ normalisation used in the previous chapters,

$$k = a^2$$

This relation is part of the structural normalisation of the repeated-root branch. In empirical calibration, one may impose this relation or treat the decay parameter and volatility scale separately, depending on the chosen normalisation and estimation strategy.

The Dirichlet $r = 1$ loading is

$H_D(\tau) = C_D\tau e^{-a^2\tau}$.

Ignoring scale, define

$h_D(\tau) = \tau e^{-a^2\tau}$.

Then

$h'_D(\tau) = e^{-a^2\tau}(1 - a^2\tau)$.

The maximum occurs when

$1 - a^2\tau = 0$,

so

$$\tau_{\max} = 1/a^2$$

Equivalently,

$a^2{=}1/\tau_{\max}$.

This gives a simple empirical diagnostic. If an estimated volatility loading has a clear interior maximum, the location of that maximum gives an estimate of the maturity decay scale under the Dirichlet repeated-root interpretation.

The same idea can be applied more generally. In Robin closure, the location and even the existence of an interior turning point depend on the anchoring parameter. But the decay scale k remains a central object. It controls the long-maturity decay of the loading.

Thus calibration should identify the decay scale explicitly rather than burying it inside an arbitrary volatility parametrisation.

13.5 Principle 4: estimate boundary anchoring from short-end behaviour

Robin closure introduces a boundary anchoring parameter.

The homogeneous Robin condition is

$\alpha H(0) + \beta H'(0) = 0.$

When $\beta \neq 0$, define $\rho = \alpha/\beta$.

Then the Robin condition may be written as

$H'(0) = -\rho H(0).$

Thus, whenever $H(0) \neq 0$,

$\rho = -H'(0)/H(0).$

This is a direct empirical interpretation. The anchoring parameter is the negative boundary slope of the volatility loading relative to its short-end level.

A low value of ρ corresponds to weak anchoring: the loading leaves the boundary relatively freely. A high value of ρ corresponds to stronger suppression of the short-end direction.

For the repeated-root Robin loading,

$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2\tau},$

we have

$H_R(0) = C_0,$

and

$H_R'(0) = C_1 - a_R^2 C_0.$

Thus, whenever $C_0 \neq 0$,

$\rho = a_R^2 - C_1/C_0.$

Equivalently,

$C_1/C_0 = a_R^2 - \rho.$

The boundary parameter can therefore be inferred from the estimated short-end level and slope of the loading, provided these quantities are observable or can be reliably extrapolated.

In practice, however, both H(0) and $H'(0)$ are extrapolated objects. The instantaneous short end is rarely observed directly as a clean maturity point, and the slope at the boundary is sensitive to curve construction, smoothing, interpolation, and the chosen parametric form. Estimates of ρ should therefore be treated as structural diagnostics rather than directly observed quantities.

This is one of the strongest empirical implications of the Robin formulation, but it is also one of the most delicate. Boundary anchoring is not merely a qualitative policy label. It becomes a parameter that can, in principle, be estimated from the shape of the maturity loading near the short end.

13.6 Principle 5: do not infer r = 1 from hump resemblance alone

The repeated-root branch is structurally important because it generates the linear-exponential component

$$\tau e^{-k\tau}$$

This component naturally produces a humped maturity profile. That profile resembles familiar humped volatility specifications used in interest-rate modelling [9, 10, 11, 13].

This resemblance is useful, but it is not proof.

A hump identifies at most a functional feature, not a root structure. Many models can generate humped or hump-like maturity profiles. Only within the IRFR admissibility calculation does the linear-exponential form arise from a repeated root of the maturity-domain operator.

A humped empirical volatility profile may therefore be consistent with the r = 1 repeated-root branch. It may motivate testing that branch. It may make the branch economically interpretable. But it does not establish that r = 1 is the correct branch.

Distinct-root branches must also be considered. They generate sums of exponentials without the linear-exponential repeated-root term, typically producing monotone or double-exponential decay rather than the pure pinned-boundary hump. Depending on coefficients and boundary conditions, such alternatives may approximate some observed shapes.

The correct empirical statement is conditional:

if the market is well described by the repeated-root IRFR branch,

then boundary selection provides a structural route to humped or partially humped volatility profiles.

The empirical task is to test this conditional claim. One should compare the repeated-root branch against distinct-root alternatives, richer factor structures, and conventional volatility specifications.

The r = 1 branch should therefore be treated as a distinguished candidate, not as an empirically established fact.

13.7 Principle 6: distinguish scale from shape

The observed magnitude of volatility and the maturity shape of volatility are different empirical objects.

In the one-factor IRFR convention, the conventional volatility coefficient is

$$\sqrt{2}aH(\tau)$$

Multiplying $H(\tau)$ by a constant changes the volatility scale, but not the underlying shape after normalisation. Conversely, changing k or ρ changes the shape of the loading, even if the total volatility level is rescaled to match observed variance.

A useful empirical decomposition is therefore:

volatility level;

maturity decay;

boundary anchoring;

factor dimension.

These are distinct dimensions of calibration.

The volatility level determines the overall size of rate shocks. The maturity-decay scale determines how quickly the loading decays at long maturities. The anchoring parameter determines how the loading behaves near the short end. The factor dimension determines how many independent stochastic sources are needed.

Without separating these objects, calibration can become misleading. A model may fit variance levels while misidentifying the boundary regime. Or it may fit a humped profile by adjusting scale and decay without capturing the correct short-end behaviour.

Thus calibration should normalise loadings before comparing shapes, and then estimate scale separately.

13.8 Principle 7: use principal components carefully

Empirical term-structure changes are often analysed using principal components.

In the one-factor IRFR setting, the covariance kernel is

$$K(\tau_1, \tau_2) = 2a^2 H(\tau_1) H(\tau_2)$$

This kernel has rank one. Its single non-zero empirical direction is proportional to $H(\tau)$.

Thus, within the one-factor restriction, the boundary-selected loading may be compared directly with the leading empirical volatility mode of forward-rate changes.

This provides a natural empirical procedure:

estimate the leading maturity loading;

normalise its scale;

compare its shape with IRFR boundary-selected forms.

However, PCA must be used carefully. Empirical yield-curve changes often require more than one factor. Level, slope, and curvature components may all be present. Measurement choices, curve construction, sampling frequency, and overlapping maturities can affect the estimated eigenfunctions.

Therefore, a mismatch between the one-factor IRFR loading and the leading empirical component does not automatically invalidate the IRFR framework. It may indicate the need for a multi-factor extension, a different sampling horizon, or a more careful mapping from observed rates to the IRFR field.

Likewise, a good match does not prove the theory. It supports the interpretation that the leading empirical maturity mode may be boundary-selected in the IRFR sense.

The empirical use of PCA should therefore be diagnostic, not definitive.

13.9 Principle 8: treat boundary regimes as estimable hypotheses

Neumann/free, Dirichlet, and Robin closures should not be assumed merely because they are mathematically available.

They should be treated as competing empirical hypotheses.

The Neumann/free benchmark predicts retained short-end stochastic loading through the ordinary exponential profile

$H_N(\tau) = A_N\, e^{-a_N^2 \tau}.$

A Dirichlet regime predicts suppressed short-end loading and, in the repeated-root branch, the pure humped component

$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau}.$

A Robin regime predicts a mixed repeated-root profile with partial short-end retention,

$H_R(\tau) = (C_0 + C_1 \tau) e^{-a_R^2 \tau}.$

These regimes imply different observable shapes near the short end. Neumann/free closure retains boundary-active stochastic exposure. Dirichlet closure suppresses it. Robin closure partially retains it while allowing a boundary-induced linear-exponential component.

Thus the empirical question is not only:

does the model fit volatility?

It is also:

which boundary law best explains the estimated maturity loading?

This suggests a regime-testing procedure. Estimate the maturity loading, normalise it, and compare the fit of Neumann/free, Dirichlet, and Robin forms. Then assess whether the inferred boundary regime is stable across time and consistent with the corresponding monetary-policy environment.

The economic interpretation matters. Dirichlet closure corresponds to strong short-end anchoring or lower-bound enforcement. Robin closure corresponds to partial policy anchoring. Neumann/free closure corresponds to the non-pinned short-end benchmark. A useful empirical result should connect the estimated boundary regime to the institutional regime it is meant to represent.

The result should be interpreted probabilistically and comparatively. The data may favour one boundary regime over another, but the conclusion remains model-dependent.

Boundary laws are structural hypotheses. They must earn their empirical interpretation through calibration and testing.

13.10 Principle 9: check positivity and admissibility separately

The previous chapters introduced lower-boundary interpretations for Dirichlet and sign-compatible Robin closures.

In the Dirichlet $r = 1$ case, the selected loading $H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau}$

is non-negative on the half-line when $C_D \geq 0$ and vanishes at $\tau = 0$. This supports a lower-boundary interpretation when the field is written relative to an admissible minimum level,

$$x(t, \tau) = x_{\min} + y(t, \tau), y(t, \tau) \geq 0$$

This lower-boundary interpretation is stronger than a static sign observation. Under the shifted multiplicative KFE structure, the Dirichlet branch admits a shifted inverse-gamma stationary density for each fixed $\tau > 0$. The same Feller classification as the positive-state benchmark makes the lower boundary y = 0 inaccessible maturity-by-maturity, provided the process starts in the positive state domain and the parameter conditions are satisfied.

The remaining qualification is field-level regularity. To lift the fixed-maturity result to the full maturity-field statement

$$y(t, \tau) \geq 0 forall \tau \geq 0,$$

one must assume or prove sufficient regularity of the field in the maturity coordinate τ. Under such regularity, fixed-τ Feller inaccessibility supports preservation of the lower-boundary state domain across the maturity field.

In the Robin case, the repeated-root loading can be written as

$H_R(\tau) = C_0[1 + (a_R^2 - \rho)\tau]e^{-a_R^2 \tau}$.

For $C_0 > 0$, non-negativity on the full half-line requires

$\rho \leq a_R^2$.

When this sign condition is satisfied and the shifted positive-state KFE structure is retained, Robin inherits the same maturity-by-maturity inverse-gamma and Feller logic. If the Robin loading changes sign, the simple lower-boundary interpretation no longer applies globally on the half-line.

Calibration should therefore separate three levels of admissibility:

loading-level sign compatibility;

fixed-maturity state-domain preservation under the shifted KFE and Feller conditions;

field-level preservation across all maturities, which requires maturity regularity.

This distinction prevents overclaiming while also recognising that the Dirichlet lower-boundary construction is supported by more than the sign of the loading alone.

13.11 Principle 10: empirical fit is evidence, not proof

The final principle is epistemic.

A successful calibration would be important. If the IRFR boundary-selected loadings fit observed volatility profiles well, if the estimated boundary regimes correspond sensibly to monetary-policy environments, and if the parameters are stable and interpretable, then the framework would gain empirical support.

But empirical fit is not proof of uniqueness or canonicality.

Other models may fit the same data. Other mechanisms may generate similar volatility shapes. The IRFR framework may be useful without being uniquely correct.

This is why the claims of the monograph must remain separated:

> admissibility is derived;
>
> minimality is argued;
>
> canonicality is conjectured;
>
> empirical relevance is tested.

Empirical work can support the last two claims. It cannot replace the mathematical derivation of admissibility, and it cannot by itself prove canonicality.

This discipline is essential if the IRFR framework is to be taken seriously.

13.12 Minimal one-factor IRFR calibration

The preceding principles suggest a minimal one-factor calibration programme.

1. Estimate the empirical maturity loading $H(\tau)$.
2. Normalise the loading to separate shape from scale.
3. Estimate k from long-maturity decay or hump location.
4. Estimate ρ from short-end slope where $H(0) \neq 0$.
5. Compare Neumann/free, Dirichlet, and Robin boundary-regime fits.
6. Check loading sign compatibility and, where relevant, lower-boundary admissibility.
7. Test stability across time and policy regimes.

This is the smallest empirical implementation that still respects the structure of the theory.

It estimates a maturity loading, separates shape from level, identifies the decay and boundary parameters, checks admissibility diagnostics, and tests the boundary regime. It does not yet require a full multi-factor extension, stochastic volatility, or option-pricing calibration.

The minimal model is therefore not the final empirical model. It is the baseline test of the structural hypothesis.

13.13 A practical calibration sequence

A more detailed empirical workflow proceeds as follows.

First, construct forward-rate or rolling-forward-rate changes from market data. The construction should be consistent with the IRFR maturity coordinate τ.

Second, estimate the empirical covariance structure across maturity. This may be done through historical changes, option-implied volatilities, or a combination of both.

Third, extract the leading maturity loading, for example through PCA or a parametric volatility fit.

Fourth, normalise the loading so that shape and scale are separated.

Fifth, estimate the decay scale k, or a^2 under the repeated-root convention. If a Dirichlet-type hump is present, the hump location gives the diagnostic $a^2 \approx 1/\tau_{\max}$.

Sixth, estimate the Robin anchoring parameter where applicable: $\rho = -H'(0)/H(0)$.

Seventh, compare boundary-regime fits: Neumann/free, Dirichlet, Robin, and, where appropriate, distinct-root alternatives.

Eighth, test stability across time and policy regimes. Parameters such as ρ may be time-varying, reflecting changing policy regimes rather than fixed structural constants.

Ninth, check admissibility diagnostics, including short-end behaviour, loading sign compatibility, and lower-boundary interpretation where relevant.

Tenth, interpret the results as evidence, not proof.

This sequence is not the only possible empirical implementation. It is a disciplined starting point.

13.14 Summary

This chapter has set out ten principles for taking the IRFR framework to data.

The central empirical object is the maturity loading $H(\tau)$. It determines how stochastic shocks are expressed across the curve. In the one-factor model, the covariance kernel is rank one and the leading maturity direction is proportional to $H(\tau)$.

The structural parameters have observable interpretations. The decay scale controls long-maturity decay and, in the Dirichlet repeated-root case, the hump location. The Robin anchoring parameter satisfies

$\rho = -H'(0)/H(0)$

when H(0) $\neq$ 0, though in practice both H(0) and $H'(0)$ may require extrapolation. Boundary regimes imply different short-end behaviours and different maturity-loading shapes.

The empirical task is to estimate these objects, compare boundary regimes, and test stability across market and policy environments.

The chapter has also sharpened the lower-boundary discipline. Dirichlet and sign-compatible Robin loadings support lower-boundary interpretations, but the empirical analyst should distinguish loading-level sign compatibility, fixed-maturity Feller support, and full field-level maturity regularity.

The central caution is that empirical resemblance is not proof. Humped volatility motivates the repeated-root branch, but does not validate it. Good calibration supports the IRFR interpretation, but does not prove uniqueness or canonicality.

The main message is:

boundary-selected loadings become empirical hypotheses when their parameters are estimated from data.

13.15 Transition

The next chapter applies this calibration discipline to the completed boundary geometry developed above. It asks whether the same estimated maturity-space structure can organise curve shape, spatial decay, maturity volatility, covariance rank, and the market-price-of-risk restriction.

The one-factor IRFR framework has been useful because it exposes the maturity-domain geometry clearly. The empirical question is now whether that geometry remains visible when confronted with fitted forward-rate data and with diagnostics that test more than cross-sectional fit.

The next chapter therefore turns from empirical principles to empirical evidence.

Chapter 14

Empirical Evidence for the Completed Boundary Geometry

14.1 Introduction

The preceding chapters completed the boundary-field representation, connected it to risk-neutral pricing, and set out the calibration discipline required for empirical use. The long-boundary construction distinguished the equilibrium long-run anchor from the realised stochastic long-maturity boundary. The resulting IRFR field is no longer described only by a short-boundary loading and an interior repeated-root deformation. It is described by a completed boundary geometry with three maturity loadings:

$$J(\tau) = \left(e^{-a^2\tau},\ \tau e^{-a^2\tau},\ 1 - e^{-a^2\tau}\right).$$

The first loading transports the short boundary into maturity space. The second loading is the repeated-root belly mode inherited from the short-boundary construction. The third loading transports the stochastic long boundary.

The purpose of the present chapter is to ask whether this completed geometry is visible in data.

This is not a test of whether historical observations can prove the boundary-field theory uniquely. They cannot. Nor is it a test of whether the theory can be made to fit a yield curve by adding enough parameters. The empirical question is narrower and more structural. If the completed boundary geometry is economically meaningful, then the same maturity-space object should organise several observables extracted from the same dataset: curve shape, the spatial decay parameter a^2, maturity volatility, covariance rank, and the admissible market-price-of-risk restriction.

The chapter therefore tests whether one estimated boundary geometry can support a linked sequence of diagnostics. Shape fit alone would show only that the representation is a useful curve model. Volatility fit would show that the same geometry has second-moment content. Covariance-rank evidence asks whether short/curvature risk and long-boundary risk are empirically separable. The market-price-of-risk test then asks whether the same boundary-risk directions also organise the risk-neutral drift restriction.

The empirical claim of this chapter is therefore cumulative. The evidence is supportive if the same estimated geometry:

1. fits the observed forward-rate curve;

2. identifies a stable long-sample spatial decay scale;
3. generates the observed maturity-volatility profile;
4. favours a two-boundary-risk covariance structure over a collapsed one-Brownian specification; and
5. provides a strong market-signal-to-noise-ratio diagnostic of the drift restriction.

The evidence reported below supports this interpretation. The benchmark specification uses U.S. fitted instantaneous forward rates at 1Y, 3Y, 5Y, 10Y, and 30Y. The 30-year fitted forward is used as the observable proxy for the stochastic long boundary. The full-sample spatial decay estimate is close to

$$a^2 \simeq 0.09.$$

The covariance evidence favours a two-boundary-risk structure separating short/curvature risk from long-boundary risk. A one-Brownian specification can reproduce some volatility ordering, but it performs materially worse on volatility scale and market-price-of-risk diagnostics. Fixed a^2 generates strong maturity-volatility geometry over structural windows, especially over rolling twenty-year samples. The market-price-of-risk restriction performs strongly against a zero-drift benchmark, while centred realised-drift tests remain unstable because realised drift is itself a fragile empirical object.

The chapter should be read as empirical support for the boundary-field structure, not as a proof of uniqueness. The tests support a specific claim: interest-rate observables contain a common boundary-field structure linking shape, volatility, covariance, and risk-neutral character.

The status of the evidence is therefore structural rather than exhaustive. This chapter is not a full econometric horse race against affine term-structure models, general HJM specifications, or market-model alternatives. It is an internal test of the linked empirical implications of the completed IRFR geometry.

Theoretical object	**Empirical diagnostic**
$J(\tau)$	curve-shape fit
a^2	spatial-decay stability
$g_S(\tau), g_X(\tau)$	maturity-volatility profile
$K = JLCL^\top J^\top$	covariance-rank test
$f(\tau) = g_S(\tau)\lambda_S + g_X(\tau)\lambda_X$	MSNR / market-price-of-risk test

14.2 Geometric Invariance and Dynamic Risk Weighting

The empirical tests in this chapter should not be read as another daily curve-fitting exercise. Their purpose is to distinguish structural geometry from dynamic market weighting. In a conventional calibration workflow, a large parameter vector may be re-solved each day, and it can become unclear which parameters describe the physical structure of the curve and which merely absorb changing risk appetite, liquidity, and positioning.

The IRFR claim is different. Conditional on a boundary regime and a spatial decay scale, the completed Jacobian is the persistent maturity-space object:

$$J(\tau) = \left(e^{-a^2\tau}, \tau e^{-a^2\tau}, 1 - e^{-a^2\tau}\right).$$

Market conditions do not rebuild this geometry from scratch each day. They change the state vector, the boundary-risk weights, the covariance scale, and the market prices of risk attached to that geometry.

This is the sense in which the theory sits between equilibrium and no-arbitrage modelling. The equilibrium side supplies an admissible maturity geometry selected by boundary closure. The no-arbitrage side applies risk-neutral pricing to the same admissible geometry through the market-price-of-risk adjustment. The moving objects are not arbitrary curve-shape coefficients; they are regime-dependent weights on a common boundary-field structure.

The empirical chain tested below is therefore:

$$\text{boundary regime} \Rightarrow J(\tau) \Rightarrow (g_S(\tau), g_X(\tau)) \Rightarrow K \Rightarrow \lambda \Rightarrow \text{pricing and risk.}$$

The geometry is the instrument; MSNR and the market price of risk describe how strongly the market plays the admissible risk directions. A successful empirical test is not simply one in which the curve is fitted. It is one in which the same geometry helps organise curve shape, maturity volatility, covariance rank, and the risk-neutral drift adjustment.

This framing also clarifies the scope of the word *unified*. If the same Jacobian-weighted covariance structure can be used for curve construction, historical covariance, no-arbitrage pricing, and risk measurement, then the model is not a separate equilibrium model, pricing model, and VaR model stitched together after the fact. It is one boundary-field geometry viewed under different measures and different risk weightings. The present chapter tests the historical shape, volatility, covariance, and MSNR layers. Swaption-cube calibration and VaR implementation are natural next tests, but they are not claimed as completed empirical results here.

14.3 From the Completed Jacobian to an Empirical State

The completed boundary-field representation writes the forward-rate field as

$$x_t(\tau) = X_t + [(r_t - X_t) + B_t\tau]\, e^{-a^2\tau}. \tag{14.1}$$

Equivalently,

$$x_t(\tau) = e^{-a^2\tau} r_t + \tau e^{-a^2\tau} B_t + \left(1 - e^{-a^2\tau}\right) X_t.$$

Thus the empirical state is

$$(r_t, B_t, X_t),$$

where r_t is the short-boundary state, B_t is the curvature or policy-deformation state, and X_t is the realised stochastic long-boundary state.

The associated maturity-loading vector is

$$J(\tau) = \left(e^{-a^2\tau},\ \tau e^{-a^2\tau},\ 1 - e^{-a^2\tau}\right). \tag{14.2}$$

For small state increments,

$$dx_t(\tau) = e^{-a^2\tau} dr_t + \tau e^{-a^2\tau} dB_t + \left(1 - e^{-a^2\tau}\right) dX_t. \tag{14.3}$$

The important empirical restriction is that the same $J(\tau)$ is used throughout the chapter. It is not re-estimated separately for shape, volatility, covariance, and market-price-of-risk tests. This is the point at which the empirical exercise becomes a test of the theory rather than a curve-fitting exercise.

14.4 Dataset and Empirical Construction

The empirical analysis uses U.S. fitted instantaneous forward rates. The benchmark maturity set is

$$\mathcal{T} = \{1, 3, 5, 10, 30\}. \tag{14.4}$$

For each date t, the observed benchmark forward-rate vector is

$$F_t = \big(f_t(1), f_t(3), f_t(5), f_t(10), f_t(30)\big). \tag{14.5}$$

The 30-year fitted forward is used as the benchmark observable proxy for the realised stochastic long boundary:

$$B_t^L := f_t(30). \tag{14.6}$$

This choice is imperfect because the true asymptotic long boundary is not directly observed. However, among the fitted forward maturities available in the benchmark dataset, the 30-year point is the closest observable representation of the long-maturity boundary.

The 20-year point is retained as a robustness check rather than included in the benchmark maturity set. It produces a similar full-sample estimate of a^2, but less stable period and rolling-window volatility correlations. The benchmark therefore uses 30Y as the long-boundary proxy, while the 20Y evidence is interpreted as part of the measurement problem associated with observing an asymptotic boundary at finite maturity.

For a given value of a^2, the observed forward curve is fitted to (14.1). The spatial decay parameter is estimated by the cross-sectional shape objective

$$\widehat{a^2} = \arg\min_{a^2} \sum_t \sum_{\tau \in \mathcal{T}} \Big[f_t(\tau) - x_t(\tau; a^2)\Big]^2. \tag{14.7}$$

The benchmark full-sample estimate is close to

$$\widehat{a^2} = 0.091893.$$

14.5 Boundary-Risk Decomposition

The completed Jacobian separates maturity geometry from stochastic rank. The empirical covariance layer asks how many risk directions are required to describe observed state variation. The benchmark specification separates short/curvature risk from long-boundary risk:

$$dr_t = \nu_r \, dW_t^S, \qquad dB_t = \nu_B \, dW_t^S, \qquad dX_t = \nu_X \, dW_t^X. \tag{14.8}$$

The Brownian directions are allowed to be correlated:

$$dW_t^S dW_t^X = \rho \, dt. \tag{14.9}$$

This gives two maturity-risk loadings:

$$g_S(\tau) = e^{-a^2\tau}\nu_r + \tau e^{-a^2\tau}\nu_B, \tag{14.10}$$

$$g_X(\tau) = \left(1 - e^{-a^2\tau}\right)\nu_X. \tag{14.11}$$

The effective maturity variance is

$$g_{\text{eff}}^2(\tau) = g_S^2(\tau) + g_X^2(\tau) + 2\rho g_S(\tau) g_X(\tau). \tag{14.12}$$

Equivalently, if

$$L = \begin{pmatrix} \nu_r & 0 \\ \nu_B & 0 \\ 0 & \nu_X \end{pmatrix}, \qquad C = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix},$$

then the model-implied maturity covariance matrix is

$$K = JLCL^\top J^\top. \tag{14.13}$$

This covariance matrix has rank at most two. It collapses to rank one only if the two boundary-risk directions are effectively the same Brownian motion. The empirical question is therefore whether the data favour a one-direction collapse or a two-boundary-risk structure.

14.6 Market Price of Risk and MSNR

The same two maturity-risk directions imply a market-price-of-risk restriction:

$$f(\tau) = g_S(\tau)\lambda_S + g_X(\tau)\lambda_X, \tag{14.14}$$

where $f(\tau)$ is the estimated physical drift of the forward-rate increment and λ_S, λ_X are the market prices of short/curvature risk and long-boundary risk.

The market prices of risk are λ_S and λ_X. The market signal-to-noise ratio, or MSNR, is not itself the market price of risk. It is the diagnostic used to evaluate how well the realised drift vector is explained by the boundary-risk projection.

Two diagnostics are reported. The zero-benchmark statistic is

$$R_{\text{zero}}^2 = 1 - \frac{\sum_\tau \left(f_\tau - \widehat{f}_\tau\right)^2}{\sum_\tau f_\tau^2}. \tag{14.15}$$

The centred statistic is

$$R_{\text{centred}}^2 = 1 - \frac{\sum_\tau \left(f_\tau - \widehat{f}_\tau\right)^2}{\sum_\tau \left(f_\tau - \bar{f}\right)^2}. \tag{14.16}$$

The zero benchmark asks whether the model explains realised drift relative to no drift. The centred benchmark asks whether the model explains the demeaned cross-maturity shape of realised drift. The centred test is stricter and less stable because realised drift is a low-signal empirical object, and because the centred denominator can be small.

14.7 Empirical Results

14.7.1 Branch selection

The branch-selection exercise shows that shape fit alone does not identify the repeated-root branch. Several branches fit the cross-section well. The decisive evidence comes from volatility transition and scale diagnostics. As the branch parameter moves toward the repeated-root case, the transition-region volatility error declines and the model/actual volatility scale ratio moves closer to one. This supports the repeated-root geometry as an empirical object, not merely as a theoretical special case.

14.7.2 One Brownian versus two boundary-risk directions

The one-Brownian model reproduces some of the ordering of maturities by volatility, but it performs materially worse on volatility scale and market-price-of-risk diagnostics. The two-boundary-risk model separates the short/curvature component from the long-boundary component and provides a stronger joint account of volatility and drift projection.

This result is important for the monograph. The completed Jacobian itself is a maturity-geometry object. It does not, by itself, force a particular empirical stochastic rank. The empirical evidence indicates that the natural two-direction implementation is more successful than a collapsed one-direction implementation.

14.7.3 Spatial decay stability

The full-sample estimate of the spatial decay parameter is close to $a^2 = 0.09$. Shorter windows show regime sensitivity, especially in the post-2020 period, but the long-sample and structural-window evidence supports a stable decay scale near this value.

This is one of the central empirical results. The same decay parameter controls the short-boundary loading, the repeated-root belly loading, and the long-boundary transmission loading. Stability of a^2 is therefore evidence for the maturity geometry itself.

14.7.4 Fixed-a^2 volatility geometry

With a^2 fixed at its full-sample estimate, the model continues to generate strong maturity-volatility geometry over structural windows. The fit is especially strong over rolling twenty-year windows. The main weak decade is 2010–2020, where zero-rate policy, quantitative easing, and policy-regime compression plausibly distort realised volatility shape.

The fixed-a^2 exercise is important because it prevents the volatility result from being a re-estimation artefact. The same decay scale estimated from shape is used to test volatility.

14.7.5 MSNR and the risk-neutral bridge

The market-price-of-risk restriction performs strongly against the zero-drift benchmark. This is the main risk-neutral diagnostic. It shows that the same two maturity-risk directions that organise covariance also provide a strong projection of realised drift.

The centred MSNR statistic is less stable. This should not be interpreted as a direct failure of the boundary-field covariance geometry. Rather, it reflects the empirical fragility of realised cross-sectional drift. The model identifies stable admissible risk-price directions; the historical realised prices attached to those directions can shift across regimes.

14.8 Interpretation

The empirical evidence supports the completed boundary-field geometry in a specific sense. It does not prove that the IRFR framework is the unique possible explanation of interest-rate data. It does not prove that risk premia are stationary. It does not prove that the 30-year fitted forward is a perfect observation of the asymptotic long boundary.

What it does show is that one boundary-field geometry can organise several objects that are usually treated separately. The same a^2 and the same completed Jacobian appear in shape, volatility, covariance, and market-price-of-risk diagnostics.

The empirical chain is:

$$a^2 \quad \Rightarrow \quad J(\tau) \quad \Rightarrow \quad (g_S, g_X) \quad \Rightarrow \quad K(\tau, \sigma) \quad \Rightarrow \quad f(\tau) = g_S(\tau)\lambda_S + g_X(\tau)\lambda_X.$$

This is the empirical unification. Shape evidence identifies the maturity geometry. Volatility evidence shows that the geometry has second-moment content. Covariance evidence supports the separation of short/curvature risk from long-boundary risk. MSNR evidence links the same risk directions to the market-price-of-risk restriction.

14.9 Limitations

Several limitations should be kept explicit.

First, the long boundary is not directly observed. The 30-year fitted forward is a proxy for the realised stochastic long boundary. This proxy is economically reasonable, but imperfect.

Second, the 20-year robustness exercise shows that finite-maturity long-boundary measurement is delicate. The 20-year point gives a similar full-sample decay estimate, but less stable volatility diagnostics. This supports the use of 30Y as the benchmark, but also confirms the difficulty of measuring an asymptotic boundary from finite data.

Third, realised drift is fragile. The zero-benchmark MSNR result is robust, but centred realised-drift tests are regime-sensitive. This is consistent with the view that historical drift is affected by monetary regimes, inflation regimes, and term-premium variation.

Fourth, the evidence supports the completed boundary geometry but does not establish uniqueness. Other models may reproduce some of the same facts. The contribution of the IRFR framework is that shape, volatility, covariance, and market-price-of-risk structure are linked by one maturity-space mechanism.

14.10 Conclusion

This chapter has tested the completed boundary-field geometry developed in the preceding chapters and disciplined by the calibration framework. The empirical evidence is supportive. Using U.S. fitted instantaneous forward rates, the benchmark specification identifies a stable long-sample decay scale near $a^2 = 0.09$, supports the 30-year point as the practical observable proxy for the stochastic long boundary, favours a two-boundary-risk covariance structure, and finds strong zero-benchmark MSNR evidence for the market-price-of-risk restriction.

The conclusion is not that the theory is proved. The conclusion is that the central empirical footprint of the theory is present: the same boundary geometry that organises the forward-rate curve also organises maturity volatility, covariance rank, and the admissible risk-neutral drift structure.

This is the empirical bridge from the completed long-boundary construction, through risk-neutral pricing and calibration discipline, to the final scope and limitations of the monograph. The theory has produced a completed Jacobian. The data suggest that this Jacobian is not merely formal. It is visible in the observable term structure.

Chapter 15

Scope, Limitations, and Further Directions

15.1 Introduction

The preceding chapters developed a minimal IRFR framework for interest-rate dynamics, completed the long-boundary geometry, connected the framework to risk-neutral pricing, set out calibration discipline, and reported empirical evidence for the completed boundary geometry.

The construction began with the instantaneous rolling forward rate,

and treated the forward curve as a stochastic field over remaining maturity. The change from calendar maturity to rolling maturity exposed the transport structure of the curve and identified the short end,

as a genuine boundary of the maturity domain.

The axioms then formalised this geometry. Axioms 1–3 defined the interior law: local stochastic evolution, rolling-maturity invariance, and equilibrium moment transport. Axiom 4 introduced boundary closure at the short end.

The Kolmogorov forward equation was then used to convert the interior axioms into calculable restrictions. This led to the affine drift-operator structure, the admissibility condition for maturity loadings, and the r-branch of admissible modes. The boundary chapters then showed how natural, Dirichlet, and Robin closure select different realised maturity loadings from the same underlying boundary-field structure.

The purpose of the present chapter is not to extend the framework beyond what has been derived. Its purpose is to state clearly what has been established, what has been supported by the construction, what remains conditional, and which questions remain open.

This distinction is important. The monograph has developed a minimal boundary-field construction. It has not proved a universal theory of all interest-rate dynamics. Nor has it established empirical validation, a complete multi-factor theory, or a complete no-arbitrage pricing framework.

The central claim is more precise:

within the one-factor IRFR field framework, admissible dynamics are constrained by rolling maturity and boundary closure.

This chapter records the scope of that claim.

15.2 What has been established

The first established result is representational.

The IRFR primitive

turns the forward curve into a field over remaining maturity. It is equivalent to the instantaneous forward curve at the level of curve values, but exposes a different geometric structure. In this representation, fixed calendar maturities roll through the maturity domain according to

This makes deterministic roll-down explicit and identifies

as the short-end boundary.

The second established result is axiomatic.

The four axioms organise the model into an interior law and a boundary law. Axioms 1–3 govern the admissible dynamics away from the boundary. Axiom 4 determines how the field is closed at the short end.

The third established result is the affine drift operator.

Under the convention

the KFE and first-moment transport imply the drift-operator form

The actual SDE drift is therefore

This separates mean-reverting displacement from the transport correction induced by rolling maturity.

The fourth established result is the r-branch.

The branch parameter is

where a^2 is the temporal drift–diffusion scale and λ^2 is the maturity-decay scale. For r $\neq$ 1, the admissible maturity modes are distinct exponentials,

and

At

the two modes coalesce. The repeated-root space contains

and

This is the algebraic source of the linear-exponential humped component.

The fifth established result is boundary selection.

The Neumann/free benchmark retains the ordinary exponential short-end stochastic loading,

Dirichlet closure suppresses short-end stochastic loading and, in the repeated-root branch, selects the boundary-induced humped mode, Robin closure gives the partially anchored mixed repeated-root loading,

Thus boundary closure does not merely impose an endpoint condition. It selects the realised maturity structure of risk.

The sixth established result is the interpretation of volatility shape.

In the one-factor model,

and the instantaneous covariance rate is

Thus the maturity loading $H(\tau)$ is the observable volatility-shape object. Boundary-selected loadings become empirical hypotheses when their parameters are estimated from data.

The structural parameters a^2, r, and, in the Robin case, ρ, are identifiable through features of the maturity loading, though only under the modelling assumptions and estimation procedures described in Chapter 11.

15.3 What has been supported, but remains conditional

Some conclusions are stronger than conjecture, but still depend on clearly stated technical conditions.

The most important example is lower-boundary admissibility in the Dirichlet branch.

In the Dirichlet lower-boundary interpretation, write

The Dirichlet boundary imposes

For $\tau > 0$, the repeated-root Dirichlet loading is

Under the positive normalisation $C_D \geq 0$, this loading is non-negative on the half-line and vanishes at the short end.

More importantly, under the shifted multiplicative KFE structure, the Dirichlet branch remains in the positive-state diffusion class. For each fixed $\tau > 0$, the shifted state may be represented in the form

with positive parameter functions. The stationary KFE then admits a shifted inverse-gamma density on

Equivalently, in the original variable,

The same Feller boundary classification as in the positive-state benchmark shows that the lower boundary y = 0 is inaccessible for each fixed $\tau > 0$, provided the process starts in the positive state domain and the parameter conditions are satisfied.

Thus the lower-boundary interpretation is not merely a static sign property of the loading. It is supported by the shifted inverse-gamma stationary density and by maturity-by-maturity Feller inaccessibility of the lower state boundary. The detailed density and boundary-classification calculation is given in Appendix B.

The remaining qualification is field-level regularity. To pass from fixed-τ positivity to the full maturity-field statement

one also requires sufficient regularity of the field in the maturity coordinate. Under such regularity,

and

are compatible with preservation of the lower-boundary state domain across the maturity field.

Robin closure inherits the same local lower-boundary logic when its loading is sign-compatible. For

with $C_0 > 0$, non-negativity on the full half-line requires

If this sign condition fails, the simple lower-boundary interpretation no longer applies globally on the half-line.

15.4 What has not been established

The framework has not proved uniqueness.

The monograph has shown that the proposed construction gives an admissible and structurally disciplined model class. It has not proved that this class is the unique possible construction satisfying the axioms, nor that no other coordinate choice or admissibility mechanism could generate similar structures.

The framework has not empirically validated the model.

The resemblance between boundary-selected loadings and familiar humped volatility forms motivates calibration and testing. It does not prove that the IRFR model is empirically correct. Nor does it prove that the r = 1 branch is the true branch. Empirical validation remains a separate task.

The framework has not developed a complete multi-factor theory.

The present monograph has deliberately focused on the one-factor case. The distinction between stochastic dimension and maturity geometry was central to the exposition. A multi-factor IRFR field may be natural, but it has not been derived here. In particular, it is not guaranteed that multi-factor extensions preserve the same admissibility structure or boundary-selection mechanism without further work.

The framework has not provided a complete derivative-pricing implementation.

Chapter 12 supplies the bond-pricing layer that was previously missing. It identifies the IRFR field with the rolling instantaneous forward-rate curve, defines zero-coupon bond prices, introduces the money-market numeraire, and imposes the HJM martingale restriction on IRFR-admissible volatility modes. Thus the conceptual risk-neutral measure for bonds is no longer left open.

What remains incomplete is the full derivative-pricing programme. Bond options, caps, floors, and swaptions require the distribution under $\mathbb{Q}$ of integrated IRFR field functionals, together with a choice of numeraire, payoff specification, factor reduction, and calibration of the market price of risk. The IRFR field value $x(t, \tau)$ is still a local rate-field coordinate rather than a strictly positive tradable price process. It cannot itself serve as a numeraire, but it can generate tradable bond prices once integrated as a forward-rate curve.

The remaining open question is therefore narrower than before: not the existence of a risk-neutral bond-pricing measure in principle, but the empirical calibration, numerical implementation, and derivative-pricing extension of the IRFR-constrained HJM framework.

These limitations are not defects. They are the boundaries of the present work.

15.5 Status of the r = 1 branch

The r = 1 branch has played a central role in the exposition because it is mathematically distinguished.

At r = 1, the maturity-decay scale and the temporal drift–diffusion scale coincide:

The two exponential modes

and

coalesce. The repeated-root space then contains the linear-exponential component

This term is humped. Under Dirichlet closure, it becomes the selected loading

This structure is close to humped volatility specifications used in interest-rate modelling. That connection is important, but it should be read carefully.

The empirical usefulness of humped volatility forms does not prove that the r = 1 branch is correct. It shows that the branch is economically interpretable and worth testing. The theoretical result is conditional:

if the repeated-root branch is relevant, boundary selection gives a structural route to humped loadings.

Whether the repeated-root branch is empirically preferred over distinct-root alternatives is a calibration question.

Thus the r = 1 branch is distinguished within the theory, not empirically established by resemblance alone.

15.6 Status of the boundary-law interpretation

The boundary-law interpretation is one of the main contributions of the monograph.

The short end is both a mathematical boundary and the location where monetary policy acts most directly. This makes boundary closure a natural place to represent economic regime.

The boundary chapters developed three canonical closures:

Neumann/free closure: short-end stochastic loading retained;

Dirichlet closure: short-end stochastic loading suppressed;

Robin closure: short-end stochastic loading partially retained.

These should be read as structural idealisations. Real monetary-policy regimes are more complex than any single boundary condition. Nevertheless, the classification is useful because it separates the interior admissible dynamics from the economic regime that selects the realised loading.

The strongest claim is therefore not that every policy regime is literally a Dirichlet or Robin boundary. The stronger and safer claim is:

policy anchoring can be represented as a boundary-selection mechanism in the IRFR field.

This is a structural interpretation. Its empirical content depends on whether estimated volatility loadings behave as the boundary-law classification predicts.

15.7 Minimality and modelling cost

The framework aims to be minimal in a specific sense.

It does not begin by specifying a volatility curve and then deriving the required arbitrage-free drift. Instead, it begins with a rolling-maturity field, imposes structural axioms, and asks which drift and loading structures are admissible.

This changes the modelling cost.

In a flexible reduced-form model, one can introduce a volatility function because it fits data. In the IRFR framework, a maturity loading must be justified as part of the admissible field structure and boundary closure.

This reduces exogenous model intervention, but it also reduces freedom. Some useful empirical shapes may be outside the minimal model. Some market features may require additional factors, stochastic volatility, time-varying boundary laws, or richer state variables.

This is the trade-off.

The minimal model is valuable not because it captures everything, but because it identifies which features can be derived from a disciplined interior-law and boundary-law construction.

This minimality also clarifies the relationship with richer nonlinear term-structure models, including quadratic specifications. Quadratic models obtain nonlinear yield behaviour by enriching the state dependence. The IRFR construction asks what nonlinear-looking maturity structure can be obtained before that step, from rolling-maturity geometry, admissibility, and boundary selection alone.

15.8 Further directions

Several further directions are natural, but they should be regarded as future work.

A first direction is empirical testing. The calibration principles of Chapter 13 provide a starting point. Actual validation requires data, estimation procedures, robustness checks, and comparison against competing models.

A second direction is to formulate the maturity-regularity conditions under which fixed-τ Feller inaccessibility lifts to lower-boundary preservation of the full maturity field. The Dirichlet branch already supports lower-boundary admissibility maturity-by-maturity; the remaining technical task is to state the appropriate function-space conditions for the whole field.

A third direction is multi-factor extension. A multi-factor IRFR model would allow several Brownian drivers and several maturity loadings. Such a model may be necessary for empirical yield-curve dynamics. However, the admissibility and boundary-selection analysis would need to be developed again, not merely assumed.

A fourth direction is time-varying boundary parameters. In the Robin case, the anchoring ratio

could vary through time. This would allow changing policy regimes to appear as changing boundary laws. The present monograph identifies this possibility but does not estimate or model such dynamics.

A fifth direction is stochastic boundary laws. Instead of fixed boundary coefficients, the boundary itself could be driven by a stochastic process representing policy uncertainty or liquidity conditions.

This may be economically realistic, but it is beyond the minimal framework.

A sixth direction is derivative-pricing implementation. Chapter 12 constructs the no-arbitrage bond-pricing layer. The next step is to price claims on bonds, swaps, caps, floors, and swaptions by analysing the risk-neutral distribution of integrated IRFR field functionals and by calibrating the Market Signal-to-Noise Ratio and market price of risk.

These directions are not part of the established result. They are consequences of the programme and possible areas for future research.

15.9 Why the minimal model still matters

The limitations just described do not undermine the value of the minimal model.

A minimal model is useful because it reveals structure that may be obscured in more flexible specifications. The one-factor IRFR framework shows how rolling maturity, moment transport, admissible modes, and boundary closure interact.

It also shows that humped maturity loading can arise as a structural consequence rather than a modelling input. In the repeated-root Dirichlet case, the hump is a boundary-selected admissible mode. In the Robin case, partial anchoring selects a mixed profile.

These are structural results. They do not require the minimal model to be the final empirical model.

The minimal model therefore plays the role of a benchmark. It gives a disciplined baseline against which extensions, empirical calibrations, and alternative specifications can be compared.

The central organising principle remains:

interior admissibility determines the mode space;

boundary closure selects the realised loading.

This principle is the main object that any extension should preserve.

15.10 Summary

This chapter has stated the scope and limitations of the IRFR construction.

The monograph has established a minimal one-factor boundary-field framework. It has shown how the IRFR primitive reveals rolling-maturity transport, how the axioms divide the theory into an interior law and a boundary law, how the KFE produces an affine drift operator, how the r-branch arises, and how boundary closure selects observable loadings.

It has also strengthened the lower-boundary interpretation of the Dirichlet branch. Under shifted positive-state dynamics, the Dirichlet branch admits a shifted inverse-gamma stationary density for fixed maturities $\tau > 0$, and the lower boundary is inaccessible maturity-by-maturity under the Feller conditions. Field-level preservation requires the corresponding regularity in the maturity coordinate.

The monograph has not proved uniqueness. It has not empirically validated the model. It has not developed a full multi-factor framework. It has not solved the no-arbitrage pricing problem associated with embedding a boundary-field short end inside a numeraire-based pricing construction.

These limitations define the boundary of the current work.

The value of the framework lies in the discipline it imposes. It shows how drift structure, maturity loading, volatility shape, boundary anchoring, and lower-boundary admissibility can be traced to specific locations in the theory.

The appropriate final claim is therefore measured:

the IRFR framework provides a minimal structural route from rolling maturity and boundary closure to admissible interest-rate dynamics.

15.11 Transition to the conclusion

The final chapter returns to the overall programme of the monograph.

It summarises the construction, restates the main results, identifies the conjectural claims, and explains why the boundary-law formulation provides a useful way to think about interest-rate dynamics.

The purpose of the conclusion is not to enlarge the claims, but to close the argument carefully: what has been derived, what has been supported by the construction, what has been motivated, and what remains open.

Chapter 16

Conclusion: A Unified Boundary-Field Theory of Interest Rates

16.1 Introduction

This monograph has developed a unified boundary-field theory of interest rates. The theory is minimal in its primitives: it begins with the instantaneous real forward rate, maturity-space dynamics, and boundary conditions. But it is unified in its consequences: the same maturity-space geometry governs curve shape, volatility, covariance, risk-neutral drift, and the natural inputs for pricing and risk measurement.

As stated, the starting point was the instantaneous rolling forward rate,

$$x(t, \tau) = F(t, t + \tau),$$

where t is calendar time and τ is remaining maturity. This change of primitive does not create a new market object. It is equivalent to the conventional instantaneous forward curve at the level of curve values. But it exposes a different geometric structure.

In rolling-maturity coordinates, fixed calendar maturities move through the maturity domain according to

$$d\tau/dt = -1$$

The forward curve therefore becomes a field over the half-line

$$\tau \in [0, \infty)\,,$$

and the short end

$$\tau = 0$$

becomes a boundary.

The central idea of the monograph is that this geometry matters. Once the forward curve is treated as a rolling maturity-domain field, drift, maturity loading, volatility shape, and policy anchoring can be organised through a common structure:

interior admissibility determines the mode space;

boundary closure selects the realised loading.

The purpose of this final chapter is to summarise what has been constructed, state the main claims at the correct strength, and identify what remains open.

16.2 The problem revisited

Interest-rate models face a persistent tension.

On the one hand, term-structure models must be flexible enough to fit observed curves, volatilities, and derivatives prices. On the other hand, excessive flexibility can obscure which features of the model are structural and which are inserted for calibration.

This tension is especially visible in volatility modelling. Humped volatility profiles are empirically useful and widely recognisable. But in many formulations, the maturity shape of volatility is specified directly. The model may fit observed behaviour, but the shape itself remains an input.

The IRFR framework asks a different question.

which maturity loadings are admissible under rolling-maturity geometry?

which of those admissible loadings are selected by boundary closure?

This changes the modelling problem. It moves part of the burden from exogenous specification to structural derivation.

The framework does not eliminate modelling choice. It relocates it. Assumptions enter through the primitive, the axioms, the admissibility calculation, and the boundary law. The point is that those assumptions are made visible.

16.3 The IRFR primitive

The first modelling commitment was to use the instantaneous rolling forward rate,

$$x(t,\tau) = F(t, t+\tau)$$

The conventional representation $F(t,T)$ labels the curve by calendar maturity. The IRFR representation labels the curve by remaining maturity.

This makes the maturity coordinate dynamic. A fixed maturity date T corresponds to

$$\tau = T - t,$$

so calendar time transports the maturity point toward the short end.

This simple observation has three consequences.

First, deterministic roll-down becomes explicit. A change in a fixed-maturity forward rate combines stochastic evolution of the field with deterministic movement along the maturity axis.

Second, the forward curve becomes a field over maturity, not merely a collection of unrelated maturity-indexed processes.

Third, the short end becomes a boundary. It is not added later as a modelling convenience. It is the endpoint of the maturity domain toward which contracts roll.

This is the geometric foundation of the monograph.

16.4 The axioms

The field representation led to four axioms.

Axiom 1 gave the field local stochastic evolution. Under the convention used in the monograph,

$$dx(t, \tau) = -f(x, t, \tau)\, dt + \sqrt{2} g(x, t, \tau)\, dW_t$$

Here f is the IRFR drift-operator coefficient, g is the diffusion-scale coefficient, and the actual SDE drift is $-f$.

Axiom 2 imposed rolling-maturity invariance. The same remaining maturity should be governed by the same structural law within a given regime, regardless of the calendar date at which it is observed.

Axiom 3 required equilibrium moment transport. The expected dynamics must be compatible with deterministic roll-down along the equilibrium curve $m(\tau)$.

Axiom 4 required boundary closure at

$$\tau = 0$$

Together, these axioms divided the model into two parts:

Axioms 1-3: interior law;

Axiom 4: boundary law.

This division is the organising principle of the framework.

16.5 The affine drift operator

The Kolmogorov forward equation turned the first three axioms into calculable restrictions.

The first-moment equation implied that the drift operator must close at the level of the first moment. Requiring compatibility with equilibrium transport led to the affine drift-operator form

$$f(x, t, \tau) = a^2[x - m(\tau)] - m'(\tau)$$

The actual SDE drift is therefore

$$-f(x,t,\tau) = -a^2[x - m(\tau)] + m'(\tau)$$

This decomposition has a clear interpretation.

$-a^2[x - m(\tau)]$

is the mean-reverting displacement term in the actual drift.

$m'(\tau)$

is the fixed-τ transport correction generated by rolling maturity.

Thus the drift structure is not merely chosen for tractability. It reflects the interaction between stochastic displacement and deterministic maturity transport.

This was the first major payoff of the IRFR formulation.

16.6 Admissible maturity modes

The diffusion-scale coefficient was then written in maturity-loading form,

$$g(\tau) = aH(\tau)$$

The KFE and moment-closure conditions imply that the maturity loading cannot be arbitrary.

The branch parameter is

$$r = \lambda^2/a^2,$$

where a^2 is the temporal drift-diffusion scale and λ^2 is the maturity-decay scale. For $r \neq 1$, the admissible maturity modes are the distinct exponentials

$e^{-a^2\tau}$ and $e^{-ra^2\tau}$.

At $r = 1$, the two modes coalesce. The repeated-root space contains $e^{-a^2\tau}$ and $\tau e^{-a^2\tau}$.

The repeated-root branch is therefore mathematically distinguished because it generates the linear-exponential component

$$\tau e^{-a^2\tau}$$

This component is humped. It vanishes at the short end, rises to an interior maximum, and decays at long maturity.

The repeated-root branch is structurally important and economically interpretable. But it is not empirically established by resemblance to humped volatility alone. It remains a distinguished branch to be tested.

16.7 Boundary selection

The boundary law determines which admissible loading is realised.

Neumann/free closure provides the non-pinned benchmark. It retains the ordinary exponential short-end stochastic loading,

$H_N(\tau) = A_N\, e^{-a_N^2 \tau}$.

The short-end loading is retained:

$H_N(0) = A_N$.

Dirichlet closure represents the idealised pinned-boundary case. It suppresses short-end stochastic loading and, in the repeated-root branch, selects the boundary-induced humped mode

$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau}$.

This loading satisfies

$H_D^{(1)}(0) = 0$.

Robin closure represents partial anchoring. It retains residual short-end stochastic loading while allowing a boundary-induced repeated-root component:

$H_R(\tau) = (C_0 + C_1 \tau) e^{-a_R^2 \tau}$.

When $C_0 \neq 0$, the same loading may be written in anchoring-ratio form as

$H_R(\tau) = C_0[1 + (a_R^2 - \rho)\tau] e^{-a_R^2 \tau}$,

where

$\rho = -H_R'(0)/H_R(0)$.

These cases demonstrate the boundary-law principle:

the interior law creates admissible modes;

the boundary law selects the realised combination.

16.8 Observable volatility

The boundary-selected maturity loading determines observable volatility shape.

In the one-factor IRFR model,

$$g(\tau) = aH(\tau)$$

Define the instantaneous covariance rate by

$$K(\tau_1, \tau_2) = (1/dt)\, \mathrm{Cov}(dx(t, \tau_1), dx(t, \tau_2))$$

Then

$$K(\tau_1, \tau_2) = 2a^2 H(\tau_1) H(\tau_2)$$

This kernel is rank one. Its maturity eigenfunction is proportional to $H(\tau)$.

Thus the empirical object corresponding to the theory is not merely a volatility level. It is a maturity loading. The shape of this loading contains information about decay, anchoring, and boundary regime.

For the Dirichlet repeated-root loading,

$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau}$,

the maximum of the loading occurs at

$$\tau_{\max} = 1/a_D^2$$

For the Robin loading,

$H_R(\tau) = (C_0 + C_1 \tau) e^{-a_R^2 \tau}$,

the anchoring parameter satisfies

$\rho = -H_R'(0)/H_R(0)$

whenever $H_R(0) \neq 0$.

These relationships turn boundary-law parameters into empirical diagnostics. But they do not by themselves validate the model. They define what should be estimated and tested.

16.9 What has been claimed

The main claim of the monograph is structural.

Within the one-factor IRFR field framework, rolling-maturity geometry and boundary closure constrain admissible interest-rate dynamics. Drift structure, maturity loading, volatility shape, and policy anchoring are not independent modelling choices. They are connected through the interior-law and boundary-law construction.

More specifically, the monograph has claimed that:

the IRFR primitive reveals transport and boundary geometry;

the axioms divide the model into interior law and boundary law;

the KFE gives calculable admissibility restrictions;

the affine drift operator follows from first-moment transport;

maturity loadings arise as admissible modes;

boundary closure selects observable loading shapes.

These are the earned claims.

The contribution is therefore not a new arbitrary volatility parametrisation, but a mechanism by which volatility loadings are selected from an admissible maturity-mode space.

The framework is also not presented as a replacement for conventional arbitrage-free representations such as Heath-Jarrow-Morton. Rather, it gives a structural route by which certain maturity loadings may arise within such representations. Heath-Jarrow-Morton describes the no-arbitrage representation of forward-rate dynamics; IRFR provides a structural route by which rolling-maturity geometry and boundary closure may select particular admissible loadings.

The pricing chapter also introduces the Market Signal-to-Noise Ratio (MSNR) as the native IRFR interpretation of the market price of risk. Under the physical measure, MSNR is the observed drift signal per unit diffusion noise. Under the risk-neutral measure, the HJM compensator replaces that drift signal. The market price of risk is therefore the correction between physical and risk-neutral MSNR.

16.10 What has not been claimed

Several stronger claims have deliberately not been made.

The monograph has not proved uniqueness. It has shown an admissible and disciplined construction, not that no other construction is possible.

It has not completed a full cross-market empirical validation of the repeated-root branch. The historical evidence reported above supports the completed boundary geometry in the benchmark U.S. dataset, but further validation is required across currencies, policy regimes, and market instruments.

It has not provided a complete multi-factor theory. Multi-factor extensions may be natural, but they require new derivation.

It has not provided a complete derivative-pricing implementation. Chapter 12 constructs a risk-neutral bond-pricing interpretation by mapping the IRFR field into zero-coupon bond prices and imposing the HJM drift restriction on admissible volatility modes. But the IRFR value $x(t,\tau)$ remains a local rate-field coordinate, not itself a strictly positive tradable price process. Pricing options on bonds, caps, floors, and swaptions still requires a tractable risk-neutral law for the relevant integrated field functionals, a payoff-specific numeraire choice, and empirical calibration of the Market Signal-to-Noise Ratio and market price of risk.

The lower-boundary claim is also scoped carefully. The monograph does not assert a fully developed field-level positivity theorem without qualification. Rather, the Dirichlet branch supports lower-boundary admissibility maturity-by-maturity through the shifted inverse-gamma stationary density and Feller inaccessibility of the lower state boundary for each fixed $\tau > 0$. The remaining technical qualification is the maturity regularity needed to lift fixed-maturity positivity to the full field statement.

These limitations define the boundary of the present contribution.

16.11 Why the framework matters

The value of the IRFR framework lies in its discipline.

It provides a way to connect objects that are often specified separately:

drift;

volatility loading;

maturity geometry;

short-end policy anchoring.

In the IRFR field, these objects are organised by a common structure.

This does not make the model less conditional. It makes the conditions clearer.

The framework also gives a structural interpretation of humped maturity loading. In the Dirichlet repeated-root case, the hump arises as a boundary-selected admissible mode. In the Robin case, partial anchoring selects a mixed profile that can retain short-end loading while still producing curvature or intermediate-maturity concentration.

Thus humped volatility can arise as a structural consequence rather than a modelling input.

The Dirichlet branch also shows how a lower-boundary regime can be supported by the shifted positive-state density and Feller boundary structure, not merely by the sign of the loading.

This is the central modelling payoff.

16.12 Why the Theory Is Unified

The word *unified* should be understood in a precise sense. The framework is not unified because it claims to solve every problem in interest-rate modelling. It is unified because the same maturity-space object is asked to carry several roles that are often separated in practice.

In the completed boundary geometry, the Jacobian determines the admissible shape through which state variables enter the forward curve. The same Jacobian determines the maturity profile of volatility once boundary-risk weights are attached. The same weighted geometry gives the covariance kernel. Under the risk-neutral measure, the same risk directions determine the market-price-of-risk adjustment and the no-arbitrage drift restriction. In a derivative-pricing implementation, the same covariance structure is the natural low-dimensional object to calibrate to caps, floors, and swaptions. Under the physical measure, the same covariance structure is the natural generator of curve shocks for risk measurement and VaR.

Thus the conceptual sequence is:

$$J(\tau) \Rightarrow \text{curve shape} \Rightarrow \text{volatility} \Rightarrow \text{covariance} \Rightarrow \text{risk-neutral pricing} \Rightarrow \text{physical risk measurement.}$$

This is the bridge between equilibrium and no-arbitrage. IRFR is an equilibrium geometry because boundary closure selects the admissible maturity structure. It is no-arbitrage compatible because the pricing chapter maps that structure into zero-coupon bond prices and applies the HJM drift restriction under the risk-neutral measure. The parameters that weight the Jacobian are allowed to be regime-dependent, but the geometry they weight is not an unconstrained daily curve-fitting object.

The present monograph has established the mathematical bridge and reported historical evidence for the shape, volatility, covariance, and MSNR layers. It has not yet completed a full swaption-cube calibration or a production VaR implementation. Those are the natural next tests of the unifying claim. If successful, they would show that the same boundary-field structure can support curve construction, option pricing, hedging, and risk measurement without introducing unrelated model layers at each stage.

16.13 Open problems

The monograph leaves several open problems.

The first is uniqueness. It remains to determine whether the admissible construction is unique under the axioms, or whether alternative admissible model classes exist.

The second is field-level lower-boundary preservation. The fixed-maturity Dirichlet argument gives shifted inverse-gamma stationary density and Feller inaccessibility for $\tau > 0$. A full maturity-field theorem requires the corresponding regularity in τ, so that the fixed-maturity result lifts to the entire half-line.

The third is extended empirical testing. The benchmark evidence reported above must be broadened to additional currencies, policy regimes, market instruments, and out-of-sample periods. Boundary regimes should be estimated, compared, and tested across policy environments.

The fourth is multi-factor extension. Observed term-structure dynamics often require more than one factor. Extending the IRFR admissibility and boundary-law construction to multiple factors is a substantial task.

The fifth is derivative-pricing implementation. Chapter 12 supplies the no-arbitrage bond-pricing layer: the IRFR field is identified with the rolling forward-rate curve, zero-coupon bond prices are defined by exponential discounting, and the HJM drift restriction imposes the martingale condition under the money-market numeraire. The remaining task is to extend this bond-pricing layer to a full derivative-pricing implementation, including risk-neutral factor dynamics, Market Signal-to-Noise Ratio calibration, and tractable pricing methods for bond options, caps, floors, and swaptions.

These problems define a research programme, not a completed theory.

16.14 Final statement

The aim of this monograph has been to develop a minimal structural route from rolling maturity and boundary closure to admissible interest-rate dynamics.

The construction remains deliberately narrow. It is a boundary-field framework with a completed benchmark geometry, not a universal model of all interest-rate dynamics. It does not claim uniqueness, and it does not claim complete empirical validation across all markets and instruments.

Its claim is that the rolling-maturity representation reveals a boundary-field structure that is hidden in the conventional calendar-maturity description. Once this structure is made explicit, the short end becomes a genuine boundary, and the modelling problem separates naturally into an interior admissibility problem and a boundary-selection problem.

Under the axioms developed here, the interior law determines the admissible mode space. Boundary closure then selects the realised maturity loading. In this way, drift, maturity loading, observable volatility shape, and policy anchoring become linked components of a single structural construction rather than independent modelling inputs.

The repeated-root branch illustrates the payoff of this construction. It produces the linear-exponential component

$$\tau e^{-a^2\tau},$$

which gives a natural humped maturity profile. Under Dirichlet closure this component is selected directly. Under Robin closure it appears as part of a mixed boundary-active loading. Thus familiar humped and partially humped volatility shapes can be interpreted as boundary-selected admissible modes.

This does not prove that markets are governed by the repeated-root branch. It does not prove that any particular boundary regime is empirically correct. It shows something more limited but still useful: a class of empirically recognisable volatility loadings can arise from a minimal boundary-field structure rather than being imposed exogenously.

The resulting theory is therefore best understood as a disciplined modelling framework. It identifies a primitive, states axioms, derives admissible dynamics, and interprets observable loading shapes through boundary closure. Its value lies not in claiming finality, but in making the structural assumptions visible and testable.

The central conclusion is:

rolling maturity creates a boundary;

the interior law determines admissible modes;

the boundary law selects observable loadings.

The forward curve is therefore not only a curve indexed by maturity dates. In rolling-maturity coordinates it is a field on a half-line, transported toward a boundary. Once that boundary is recognised, volatility shape can be interpreted not merely as a calibration input, but as the observable footprint of boundary selection.

This is the minimal boundary-field theory of interest rates developed in this monograph.

Appendix A

Notation and Conventions

A.1 Purpose of this appendix

This appendix collects the notation and sign conventions used throughout the monograph. The main distinctions are between calendar maturity and remaining maturity, between the drift-operator coefficient and the actual SDE drift, and between the diffusion-scale coefficient and the conventional volatility coefficient.

A.2 Time and maturity variables

Calendar time is denoted by

$$t.$$

Calendar maturity is denoted by

$$T.$$

Remaining maturity is denoted by

$$\tau.$$

The relationship is

$$\tau = T - t, \qquad \frac{d\tau}{dt} = -1.$$

A.3 Forward rate and IRFR primitive

The conventional instantaneous forward rate is

$$F(t, T).$$

The instantaneous rolling forward rate is

$$x(t, \tau) = F(t, t + \tau),$$

with inverse relation

$$F(t, T) = x(t, T - t).$$

Thus the IRFR is not a new market object at the level of curve values. It is the conventional forward curve expressed in rolling-maturity coordinates.

A.4 Transport identity

Differentiating

$$F(t,T) = x(t, T-t)$$

along a fixed calendar maturity gives

$$\frac{d}{dt}F(t,T) = \partial_t x(t,\tau) - \partial_\tau x(t,\tau), \qquad \tau = T - t.$$

The first term is calendar-time field evolution. The second is deterministic maturity transport.

A.5 IRFR stochastic convention

The one-factor IRFR field is written as

$$dx(t,\tau) = -f(x,t,\tau)\,dt + \sqrt{2}\,g(x,t,\tau)\,dW_t.$$

Here f is the IRFR drift-operator coefficient. The actual SDE drift is

$$-f(x,t,\tau).$$

The diffusion-scale coefficient is g, and the conventional SDE volatility coefficient is

$$\sqrt{2}\,g(x,t,\tau).$$

The factor $\sqrt{2}$ is a normalisation convention.

A.6 Kolmogorov forward equation convention

Under

$$dx = -f\,dt + \sqrt{2}\,g\,dW_t,$$

the density $p(x,t;\tau)$ satisfies

$$\frac{\partial p}{\partial t} = \frac{\partial}{\partial x}(fp) + \frac{\partial^2}{\partial x^2}(g^2 p).$$

Equivalently, if

$$v = 2g^2,$$

then

$$\frac{\partial p}{\partial t} = \frac{\partial}{\partial x}(fp) + \frac{1}{2}\frac{\partial^2}{\partial x^2}(vp).$$

The first-moment equation is

$$\frac{\partial M_1}{\partial t} = -\mathbb{E}[f(x,t,\tau)].$$

A.7 Affine drift operator

The equilibrium or reference curve is denoted by $m(\tau)$. The affine drift-operator coefficient derived in the monograph is

$$f(x,t,\tau) = a^2[x - m(\tau)] - m'(\tau).$$

Therefore the actual SDE drift is

$$-f(x,t,\tau) = -a^2[x - m(\tau)] + m'(\tau).$$

The first term is mean-reverting displacement. The second is the fixed-τ transport correction.

A.8 Diffusion loading and covariance rate

In the one-factor loading representation,

$$g(\tau) = aH(\tau).$$

The conventional volatility loading is

$$\sqrt{2}\, aH(\tau).$$

The instantaneous covariance of increments is

$$\operatorname{Cov}\big(dx(t,\tau_1), dx(t,\tau_2)\big) = 2a^2 H(\tau_1)H(\tau_2)\, dt.$$

The covariance-rate kernel is

$$K(\tau_1,\tau_2) = \frac{1}{dt}\operatorname{Cov}\big(dx(t,\tau_1), dx(t,\tau_2)\big) = 2a^2 H(\tau_1)H(\tau_2).$$

A.9 Admissible maturity modes

In the distinct-root case,

$$H(\tau) = C_1 e^{-k_1\tau} + C_2 e^{-k_2\tau}.$$

In the repeated-root case,

$$H(\tau) = (C_1 + C_2\tau)e^{-k\tau}.$$

Under the central $r = 1$ normalisation,

$$k = a^2.$$

The repeated-root branch contains the linear-exponential component

$$\tau e^{-k\tau},$$

which is the source of the humped loading selected by Dirichlet closure.

A.10 The r-branch

The branch parameter is

$$r = \frac{\lambda^2}{a^2},$$

where λ^2 is the maturity-decay scale and a^2 is the temporal drift-diffusion scale. The case $r \neq 1$ gives distinct exponential modes. The case $r = 1$ gives coalescence and the repeated-root loading.

A.11 Boundary regimes

The maturity domain is

$$\tau \in [0, \infty)\,,$$

with boundary at $\tau = 0$. The monograph studies three canonical boundary regimes.

The Neumann/free benchmark retains the ordinary exponential short-end stochastic mode:

$$H_N(\tau) = A_N\, e^{-a_N^2\tau}.$$

The terminology follows the convention fixed in Chapter 6: the operative condition is non-pinning, represented here by the retained exponential loading.

Dirichlet closure suppresses short-end stochastic loading and, in the repeated-root branch, selects

$$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau}.$$

Robin closure represents partial anchoring and selects the mixed repeated-root loading

$$H_R(\tau) = (C_0 + C_1 \tau) e^{-a_R^2 \tau}.$$

When $C_0 \neq 0$, this can be written as

$$H_R(\tau) = C_0 \big[1 + (a_R^2 - \rho)\tau\big] e^{-a_R^2 \tau},$$

where

$$\rho = -\frac{H_R'(0)}{H_R(0)} = a_R^2 - \frac{C_1}{C_0}.$$

A.12 Lower-boundary convention

For the Dirichlet lower-boundary interpretation, write

$$y(t,\tau) = x(t,\tau) - x_{\min}.$$

The Dirichlet boundary imposes

$$y(t,0) = 0.$$

For $\tau > 0$, the shifted positive-state structure supports a shifted inverse-gamma stationary density and Feller inaccessibility of the lower state boundary maturity-by-maturity. Full field-level lower-boundary preservation requires regularity in the maturity coordinate.

Appendix B

Algebra of Affine Closure, the r-Branch, and Boundary-Selected Loadings

B.1 Purpose of this appendix

This appendix records the algebra behind Chapters 4–9: the KFE convention, affine drift closure, the r-branch, the coalescing-root Dirichlet limit, and the boundary-selected loadings.

B.2 Stochastic convention and KFE

For fixed remaining maturity τ, write

$$dx(t,\tau) = -f(x,t,\tau)\,dt + \sqrt{2}\,g(x,t,\tau)\,dW_t.$$

The Kolmogorov forward equation is

$$\frac{\partial p}{\partial t} = \frac{\partial}{\partial x}(fp) + \frac{\partial^2}{\partial x^2}(g^2 p).$$

B.3 First moment

Let

$$M_1(t,\tau) = \mathbb{E}[x(t,\tau)] = \int xp(x,t;\tau)\,dx.$$

Then

$$\frac{\partial M_1}{\partial t} = \int x\frac{\partial p}{\partial t}\,dx.$$

Substituting the KFE and integrating by parts gives

$$\int x\frac{\partial}{\partial x}(fp)\,dx = -\int fp\,dx = -\mathbb{E}[f],$$

and

$$\int x\frac{\partial^2}{\partial x^2}(g^2 p)\,dx = 0.$$

Therefore

$$\frac{\partial M_1}{\partial t} = -\mathbb{E}[f(x,t,\tau)].$$

B.4 Affine closure

First-moment closure requires $\mathbb{E}[f]$ to depend only on M_1. Thus take

$$f(x,t,\tau) = A(\tau)x + B(\tau).$$

Choosing $A(\tau) = a^2$, imposing mean reversion toward $m(\tau)$, and incorporating transport gives

$$B(\tau) = -a^2 m(\tau) - m'(\tau).$$

Thus

$$f(x,t,\tau) = a^2[x - m(\tau)] - m'(\tau),$$

and the actual drift is

$$-f(x,t,\tau) = -a^2[x - m(\tau)] + m'(\tau).$$

B.5 Diffusion loading

The one-factor diffusion scale is

$$g(\tau) = aH(\tau).$$

Hence

$$\mathrm{Cov}\big(dx(t,\tau_1), dx(t,\tau_2)\big) = 2a^2 H(\tau_1)H(\tau_2)\,dt,$$

and

$$K(\tau_1,\tau_2) = 2a^2 H(\tau_1)H(\tau_2).$$

B.6 The r-branch

Let a^2 be the temporal drift-diffusion scale and λ^2 the maturity-decay scale. Define

$$r = \frac{\lambda^2}{a^2} > 0.$$

For $r \neq 1$, the admissible modes are

$$e^{-a^2\tau}, \qquad e^{-ra^2\tau}.$$

Writing $k = a^2$, these are $e^{-k\tau}$ and $e^{-rk\tau}$. At $r = 1$, the roots coalesce and the repeated-root space is spanned by

$$e^{-k\tau}, \qquad \tau e^{-k\tau}.$$

B.7 Distinct-root Dirichlet loading

For $r_D \neq 1$, the distinct-root Dirichlet loading is

$$H_D^{(r)}(\tau) = \frac{r_D}{1+r_D}(x_D - x_\infty)\left(e^{-r_D a_D^2 \tau} - e^{-a_D^2 \tau}\right).$$

The corresponding variance rate is

$$\frac{\partial}{\partial t}\operatorname{Var}(x_{t,t+\tau}^D) = 2a_D^2\left[H_D^{(r)}(\tau)\right]^2.$$

Thus

$$\frac{\partial}{\partial t}\operatorname{Var}(x_{t,t+\tau}^D) = 2a_D^2\left(\frac{r_D}{1+r_D}\right)^2(x_D - x_\infty)^2\left(e^{-r_D a_D^2 \tau} - e^{-a_D^2 \tau}\right)^2.$$

B.8 The repeated-root Dirichlet limit

The case $r_D = 1$ is a coalescing-root case. Consider

$$\frac{e^{-r_D a_D^2 \tau} - e^{-a_D^2 \tau}}{1 - r_D}.$$

As $r_D \to 1$,

$$\frac{e^{-r_D a_D^2 \tau} - e^{-a_D^2 \tau}}{1 - r_D} \longrightarrow a_D^2 \tau e^{-a_D^2 \tau}.$$

Therefore the non-trivial repeated-root Dirichlet loading is

$$H_D^{(1)}(\tau) = \bar{C}_D \tau e^{-a_D^2 \tau}.$$

The scalar coefficient is a normalisation convention: the limiting factor a_D^2 and any sign inherited from the orientation of the distinct-root basis are absorbed into $\bar{C}_D$. The boundary-selected object is the repeated-root mode shape $\tau e^{-a_D^2 \tau}$, not a separate sign convention.

B.9 Dirichlet hump and variance rate

For

$$H_D^{(1)}(\tau) = C_D \tau e^{-a_D^2 \tau},$$

the variance rate is

$$\frac{\partial}{\partial t}\operatorname{Var}(x_{t,t+\tau}^D) = 2a_D^2 C_D^2 \tau^2 e^{-2a_D^2 \tau}.$$

Ignoring scale, let

$$h_D(\tau) = \tau e^{-a_D^2 \tau}.$$

Then

$$h_D'(\tau) = e^{-a_D^2 \tau}(1 - a_D^2 \tau),$$

so the maximum occurs at

$$\tau_{\max} = \frac{1}{a_D^2}.$$

B.10 Shifted inverse-gamma density

Write the field relative to a lower admissible level:

$$y(t,\tau) = x(t,\tau) - x_{\min}, \qquad y(t,\tau) \geq 0.$$

For each fixed $\tau > 0$, suppose the shifted Dirichlet dynamics remain in the multiplicative positive-state class

$$dy_t = \kappa_D(\tau)\big(\theta_D(\tau) - y_t\big)\,dt + \sqrt{2}\,\eta_D(\tau)y_t\,dW_t.$$

The stationary KFE admits the inverse-gamma density

$$p_D(y;\tau) = \frac{\beta_D(\tau)^{\alpha_D(\tau)}}{\Gamma(\alpha_D(\tau))} y^{-\alpha_D(\tau)-1} \exp\left(-\frac{\beta_D(\tau)}{y}\right), \qquad y > 0,$$

where

$$\alpha_D(\tau) = 1 + \frac{\kappa_D(\tau)}{\eta_D^2(\tau)}, \qquad \beta_D(\tau) = \frac{\kappa_D(\tau)\theta_D(\tau)}{\eta_D^2(\tau)}.$$

Equivalently, in the original variable,

$$x > x_{\min}.$$

This is a multiplicative positive-state diffusion, not a CIR square-root diffusion; the CIR Feller condition does not apply.

B.11 Feller inaccessibility

For fixed $\tau > 0$, the diffusion coefficient is

$$\sigma_D(y,\tau) = \sqrt{2}\,\eta_D(\tau)y,$$

so $\sigma_D(0,\tau) = 0$. The drift at the lower boundary is

$$b_D(0,\tau) = \kappa_D(\tau)\theta_D(\tau) > 0.$$

Thus the lower boundary $y = 0$ is inaccessible for each fixed $\tau > 0$, provided the process starts in the positive state domain and the positive-state parameter conditions hold. To lift this to

$$y(t,\tau) \geq 0 \qquad \text{for all } \tau \geq 0,$$

one also requires sufficient regularity in τ. At $\tau = 0$, Dirichlet closure imposes

$$y(t,0) = 0.$$

B.12 Neumann/free benchmark

The Neumann/free benchmark retains the ordinary exponential stochastic mode:

$$H_N(\tau) = A_N\,e^{-a_N^2\tau}.$$

The short-end loading is

$$H_N(0) = A_N.$$

The variance rate is

$$\frac{\partial}{\partial t}\operatorname{Var}(x_{t,t+\tau}^N) = 2a_N^2 A_N^2 e^{-2a_N^2\tau}.$$

Thus stochastic exposure is retained at the short end whenever $A_N \neq 0$. This is the non-pinned benchmark against which Dirichlet and Robin anchoring are compared.

B.13 Robin loading and anchoring parameter

The Robin repeated-root loading is

$$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2\tau}.$$

At the boundary,

$$H_R(0) = C_0, \qquad H_R'(0) = C_1 - a_R^2 C_0.$$

The anchoring parameter is

$$\rho = -\frac{H_R'(0)}{H_R(0)} = a_R^2 - \frac{C_1}{C_0},$$

whenever $C_0 \neq 0$. Hence

$$H_R(\tau) = C_0\big[1 + (a_R^2 - \rho)\tau\big]e^{-a_R^2\tau}.$$

Writing $k = a_R^2$, this becomes

$$H_R(\tau) = C_0[1 + (k - \rho)\tau]e^{-k\tau}.$$

B.14 Robin sign compatibility

For $C_0 > 0$, the Robin loading

$$H_R(\tau) = C_0[1 + (a_R^2 - \rho)\tau]e^{-a_R^2\tau}$$

is non-negative on the whole half-line if and only if

$$\rho \leq a_R^2.$$

If $\rho > a_R^2$, the loading changes sign at sufficiently long maturity and no longer supports a simple globally non-negative lower-boundary interpretation.

B.15 Summary

The appendix establishes the algebraic spine of the monograph:

$$dx = -f\,dt + \sqrt{2}g\,dW_t,$$
$$f(x,t,\tau) = a^2[x - m(\tau)] - m'(\tau),$$
$$g(\tau) = aH(\tau),$$
$$r = \frac{\lambda^2}{a^2}.$$

For $r = 1$, the repeated-root mode space contains

$$e^{-a^2\tau}, \qquad \tau e^{-a^2\tau}.$$

Boundary closure selects the realised loading:

$$H_N(\tau) = A_N\,e^{-a_N^2\tau},$$
$$H_D^{(1)}(\tau) = C_D\tau e^{-a_D^2\tau},$$
$$H_R(\tau) = (C_0 + C_1\tau)e^{-a_R^2\tau}.$$

The humped Dirichlet loading is therefore not imposed as a primitive volatility input. It is the boundary-selected coalescing-root component of the admissible IRFR mode space.

References

[1] Ahn, D.-H., Dittmar, R. F., and Gallant, A. R. (2002). Quadratic term structure models: Theory and evidence. *The Review of Financial Studies*, 15(1), 243–288.

[2] Brace, A., Gatarek, D., and Musiela, M. (1997). The market model of interest rate dynamics. *Mathematical Finance*, 7(2), 127–147.

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

[4] Duffie, D. and Kan, R. (1996). A yield-factor model of interest rates. *Mathematical Finance*, 6(4), 379–406.

[5] Filipović, D. (2001). *Consistency Problems for Heath-Jarrow-Morton Interest Rate Models*. Lecture Notes in Mathematics, Vol. 1760. Springer.

[6] 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(1), 77–105.

[7] Hull, J. and White, A. (1990). Pricing interest-rate-derivative securities. *The Review of Financial Studies*, 3(4), 573–592.

[8] Leippold, M. and Wu, L. (2002). Asset pricing under the quadratic class. *Journal of Financial and Quantitative Analysis*, 37(2), 271–295.

[9] Mercurio, F. and Moraleda, J. M. (2000). An analytically tractable interest rate model with humped volatility. *European Journal of Operational Research*, 120(1), 205–214.

[10] Mercurio, F. and Moraleda, J. M. (2001). A family of humped volatility models. *The European Journal of Finance*, 7(2), 93–116.

[11] Moraleda, J. M. and Vorst, T. C. F. (1997). Pricing American interest rate claims with humped volatility models. *Journal of Banking & Finance*, 21(8), 1131–1157.

[12] Musiela, M. and Rutkowski, M. (2005). *Martingale Methods in Financial Modelling*. 2nd ed. Springer.

[13] Ritchken, P. and Chuang, I. (1999). Interest rate option pricing with volatility humps. *Review of Derivatives Research*, 3, 237–262.

[14] Ritchken, P. and Sankarasubramanian, L. (1995). Volatility structures of forward rates and the dynamics of the term structure. *Mathematical Finance*, 5(1), 55–72.

[15] Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2), 177–188.