Volume II

Financial Wave Mechanics

Fourier Methods · Finite Differences & Elements · Monte Carlo · Calibration by Asset Class · Real-Time Systems

Numerical Recipes
for Financial Wave Mechanics

Abdelkader Bousabaa

University of Paris-Saclay

Volume II — Second Edition — 2026

Seven Lions Editions

Contents

Preface xix

I Analytical and Semi-Analytical Methods 1

1 Fourier Methods and the Quantum Characteristic Function 3
1.1 The Quantum Characteristic Function 4
1.1.1 Definition and decomposition 4
1.1.2 Individual components 5
1.1.3 The cross characteristic function 5
1.1.4 Complete closed-form expression 5
1.2 FFT-Based Option Pricing: The Quantum Carr–Madan Formula . 7
1.2.1 From the characteristic function to option prices 7
1.2.2 Evaluating ϕ_Q at complex frequency 8
1.2.3 FFT discretisation 8
1.2.4 Choice of α and truncation error 10
1.2.5 The Lewis–Lipton formula: a damping-free alternative . . . 11
1.3 The COS Method for Quantum Densities 12
1.3.1 Principle 12
1.3.2 Pricing formula 13
1.3.3 Truncation interval for quantum densities 13
1.3.4 Convergence 13
1.3.5 COS vs FFT: when to use which 14
1.4 Worked Example: Quantum Smile via FFT 14
1.4.1 Parameter set 14
1.4.2 Python implementation 15
1.4.3 Comparison: FFT vs direct integration 17
1.4.4 Structure of the smile 17
1.5 Extensions and Complements 18
1.5.1 Multi-maturity pricing 18
1.5.2 Greeks via Fourier differentiation 18
1.5.3 Density reconstruction via IFFT 18
1.5.4 Implied volatility extraction 19
1.5.5 Fractional FFT: decoupling η and λ 20
1.5.6 Application to credit portfolio loss distributions 20

1.5.7 Relation to the closed-form Vol. I formula 21
1.6 Chapter Summary 22

2 Perturbation Theory and the WKB Approximation 25
2.1 Perturbation Expansion in $\hbar_f$ 26
2.1.1 The quantum parameter 26
2.1.2 Quadratic variation and the origin of ε 26
2.1.3 Leading order: the packet-separation channel 27
2.1.4 The dispersion channel: the quantum vega adjustment . . . 27
2.1.5 Second-order correction and the volga term 28
2.1.6 Correction to the interference term 29
2.1.7 Accuracy table 29
2.1.8 Python implementation 30
2.2 The WKB Approximation for Tunnelling 33
2.2.1 Stationary financial Schrödinger equation 33
2.2.2 WKB ansatz and the eikonal equation 33
2.2.3 Tunnelling amplitude 33
2.2.4 Application: down-and-in barrier option 34
2.2.5 Application: quantum default probability 35
2.3 First Passage Times: Exact Results and WKB 35
2.3.1 The inverse Gaussian density 35
2.3.2 Laplace transform and the WKB exponent 36
2.3.3 Saddle-point path and the inverse Gaussian exponent . . . 36
2.4 Saddle-Point Approximation for Path Integrals 38
2.4.1 Path integral representation of the propagator 38
2.4.2 The classical (dominant) path 38
2.4.3 One-loop quantum correction 39
2.4.4 Application: long-maturity barrier options 39
2.5 WKB at the Turning Point: Connection Formulas and Fresnel Integrals 40
2.5.1 Failure of the WKB at turning points 40
2.5.2 Linear turning point and the Airy equation 40
2.5.3 Fresnel integrals in the transition region 41
2.5.4 Uniform WKB approximation for ATM options 41
2.6 Validity Domains and Method Selection 42
2.6.1 Criteria 42
2.6.2 Decision diagram 43
2.6.3 Accuracy of perturbation theory vs. ε 43
2.7 Chapter Summary 43

II Grid-Based Methods 47

3 Finite Difference Methods for the Financial Schrödinger Equation 49
3.1 Discretisation of the Financial Schrödinger Equation 50

3.1.1 The equation . . . 50
3.1.2 Grid . . . 50
3.1.3 Finite-difference operators . . . 51
3.2 Explicit Scheme . . . 52
3.2.1 Forward Euler discretisation . . . 52
3.2.2 Von Neumann stability analysis . . . 52
3.2.3 Stability with a small imaginary potential . . . 53
3.3 Crank–Nicolson Scheme . . . 53
3.3.1 Cayley's form and unitarity . . . 53
3.3.2 Scalar form . . . 54
3.3.3 Von Neumann analysis of Crank–Nicolson . . . 54
3.3.4 The Thomas algorithm . . . 54
3.3.5 Algorithm summary . . . 55
3.4 The Black–Scholes Finite Difference Scheme . . . 56
3.4.1 The Black–Scholes operator and implicit scheme . . . 56
3.4.2 The M-matrix property and discrete maximum principle . . . 57
3.4.3 The Crank–Nicolson scheme for Black–Scholes . . . 58
3.5 Convergence Analysis and Grid Design . . . 59
3.5.1 Convergence near the payoff singularity . . . 59
3.5.2 Log-variable transformation . . . 59
3.5.3 Non-uniform grids . . . 60
3.6 Boundary Conditions for Financial Problems . . . 61
3.6.1 Hard walls (Dirichlet boundaries) . . . 61
3.6.2 Reflecting boundaries (natural boundaries) . . . 62
3.6.3 Absorbing and transparent boundaries (open domain) . . . 62
3.6.4 Boundary conditions in the tridiagonal system . . . 63
3.7 American Options: The Obstacle Problem . . . 63
3.7.1 The linear complementarity formulation . . . 63
3.7.2 Discrete variational inequality . . . 64
3.7.3 PSOR algorithm . . . 64
3.8 Numerical Dispersion and Phase Accuracy . . . 65
3.8.1 Phase velocity of the continuous FSE . . . 65
3.8.2 Numerical phase velocity of the CN scheme . . . 66
3.8.3 Resolution criterion . . . 66
3.8.4 Comparison with Black–Scholes finite differences . . . 66
3.9 Worked Example: The Calibrated Two-Signal State . . . 67
3.9.1 Setup . . . 67
3.9.2 Potential and initial condition . . . 68
3.9.3 Grid and convergence . . . 68
3.9.4 Convergence rate . . . 69
3.9.5 Python implementation . . . 70
3.9.6 What happens when a wall is added . . . 71
3.9.7 Comparison of schemes . . . 71
3.10 Chapter Summary . . . 72

4 Finite Element Methods **75**
4.1 Weak Formulation of the One-Dimensional Model Problem 76
4.1.1 The model boundary-value problem 76
4.1.2 Sobolev spaces 76
4.1.3 Variational formulation and well-posedness 77
4.2 Galerkin Discretisation 78
4.2.1 The Galerkin problem 78
4.2.2 Céa's lemma 78
4.2.3 P1 hat functions and interpolation error 79
4.2.4 Galerkin–least-squares stabilisation 80
4.3 P1 Finite Elements for the Black–Scholes Equation 81
4.3.1 Weak formulation of the parabolic BS problem 81
4.3.2 Implicit Euler time discretisation 82
4.3.3 Stability estimates 82
4.3.4 The mass matrix 83
4.3.5 The stiffness matrix 84
4.4 FEM–Crank–Nicolson for the Financial Schrödinger Equation . . . 85
4.4.1 Sesquilinear weak formulation 85
4.4.2 Mass matrix, stiffness matrix, and the FEM–CN system . . 85
4.4.3 Unitarity of the FEM–CN propagator 86
4.4.4 Worked example: the calibrated two-signal state on a graded mesh 86
4.5 Non-Uniform Meshes and A Posteriori Error Estimation 89
4.5.1 The key advantage: mesh-free stiffness assembly 89
4.5.2 Mesh refinement near the strike 89
4.5.3 A posteriori error estimator 90
4.6 Two-Asset Options and 2D Finite Elements 90
4.6.1 The two-asset Black–Scholes PDE 90
4.6.2 Triangulations of Ω 91
4.6.3 Nodal basis (tent functions) 91
4.6.4 Barrier options on curved domains 92
4.6.5 Intermittent barriers 93
4.6.6 Comparison: FEM versus FD for two-asset problems 94
4.7 Chapter Summary 94

5 Spectral Methods **97**
5.1 Hermite Function Basis 97
5.1.1 Hermite Polynomials and Functions 97
5.1.2 Eigenvalue Problem for the Financial Harmonic Oscillator . 98
5.1.3 Recurrence Relations and Numerical Evaluation 99
5.1.4 Exponential Convergence for Analytic Wave Functions . . . 99
5.1.5 Scale Matching: When a Gaussian Really Is One Mode . . 100
5.2 Chebyshev and Fourier Spectral Methods 101
5.2.1 Chebyshev Polynomials: Definition and Key Formulas . . . 101

5.2.2 Chebyshev Approximation on a Bounded Domain 101
5.2.3 Derivatives and Integrals of a Chebyshev Fit 103
5.2.4 Three Approaches to Spectral Discretisation 103
5.2.5 Fourier Spectral Method for Periodic Problems 104
5.2.6 Differentiation Matrices 105
5.2.7 Spectral and Physical-Space Representations 105
5.2.8 Variable Coefficients and Local-Volatility Models 106
5.3 A Worked Spectral Example 107
5.3.1 Problem Setup 107
5.3.2 Collocation Nodes and Differentiation Matrix 108
5.3.3 Assembling and Solving the System 108
5.3.4 Convergence: Numerical Results 108
5.3.5 The Calibrated Two-Signal State 109
5.4 Convergence Theory 111
5.4.1 Spectral vs. Algebraic Convergence 111
5.4.2 Gibbs Phenomenon and Filtering 112
5.4.3 Optimal Basis Choice as a Function of Wave-Function Smoothness 112
5.5 Non-Self-Adjoint Hamiltonians 113
5.5.1 Complex Potential and Absorption 113
5.5.2 Biorthogonal Expansion 114
5.5.3 Application: Knock-Out Options with a Spectral Method . 114
5.6 Tensor Product Spectral Methods for Multi-Asset 115
5.6.1 The Tensor Product Construction 115
5.6.2 Computational Cost and Practical Limits 115
5.6.3 Comparison with Quasi-Monte Carlo for $d = 2, 3$ 116
5.7 Chapter Summary 117

III Monte Carlo Methods 121

6 Classical and Quantum Monte Carlo 123
6.1 Brownian Path Sampling 123
6.1.1 Generating Standard Normal Deviates 123
6.1.2 Euler–Maruyama Discretisation 124
6.1.3 Brownian Bridge Construction 125
6.1.4 Correlated Multi-Asset Paths via Cholesky Decomposition . 126
6.1.5 Ornstein-Uhlenbeck Process Simulation 127
6.2 Quantum Importance Sampling 128
6.2.1 From Path Integral to Monte Carlo Estimator 128
6.2.2 The QIS Estimator 128
6.2.3 The Sign Problem 129
6.2.4 Financial Mitigation Strategies 131
6.3 Variance Reduction Techniques 132
6.3.1 Control Variates 132

6.3.2 Optimal Control Variates via the Martingale Representation Theorem . . . 133
6.3.3 Antithetic Sampling . . . 133
6.3.4 Stratified Sampling . . . 134
6.3.5 Importance Sampling for Deep OTM Options . . . 134
6.3.6 Adaptive Monte Carlo: VEGAS and MISER . . . 135
6.4 Convergence Diagnostics . . . 136
6.4.1 Theoretical Foundations: LLN, CLT, and Slutsky's Theorem 136
6.4.2 Practical Standard Error Estimation . . . 137
6.4.3 Effective Sample Size for Importance-Weighted Estimators 138
6.4.4 Convergence Monitoring . . . 138
6.5 Worked Examples: Basket Pricing and the Quantum Sweep . . . 138
6.5.1 Five-Asset Basket, Classical Monte Carlo . . . 139
6.5.2 Quantum Sweep at the Calibrated $\hbar_f$ . . . 140
6.5.3 Python Implementation . . . 143
6.6 Chapter Summary . . . 144

7 Quasi-Monte Carlo and Low-Discrepancy Sequences 147
7.1 Discrepancy and the Koksma–Hlawka Inequality . . . 147
7.1.1 Equidistributed Sequences . . . 147
7.1.2 Star Discrepancy . . . 148
7.1.3 The Koksma–Hlawka Inequality . . . 149
7.1.4 Discontinuous Payoffs and the $N^{-2/3}$ Rate . . . 150
7.2 Low-Discrepancy Sequences . . . 150
7.2.1 Van der Corput Sequence . . . 150
7.2.2 The Halton Sequence . . . 151
7.2.3 Latin Hypercube Sampling . . . 152
7.3 Sobol' Sequences . . . 152
7.3.1 Construction via Direction Numbers . . . 152
7.3.2 Gray Code Efficiency . . . 153
7.3.3 Scrambled Sobol' and Owen Randomisation . . . 153
7.3.4 Numerical Benchmark: Smooth vs. Hard-Boundary Integrands 154
7.3.5 Effective Dimension . . . 156
7.4 Brownian Bridge Construction for QMC . . . 156
7.4.1 Exact Brownian Path Simulation . . . 156
7.4.2 The Brownian Bridge . . . 156
7.4.3 Why the Bridge Dramatically Reduces Effective Dimension 157
7.4.4 PCA Path Construction . . . 158
7.5 Worked Example: QMC vs. MC for a Quantum Barrier Option . . 158
7.5.1 Setup . . . 159
7.5.2 Sequential MC vs. QMC (Sobol' + Bridge) . . . 159
7.5.3 Python Implementation . . . 160
7.5.4 Quantum Barrier Option . . . 162
7.6 Chapter Summary . . . 163

8 MCMC and Bayesian Calibration **167**
8.1 Market Data Preparation: the Implied Volatility Bootstrap 167
8.1.1 From Observed Prices to Implied Volatilities 168
8.1.2 Measurement Uncertainty from Bid–Ask Spreads 168
8.1.3 The Volatility Smile as Calibration Input 169
8.1.4 S&P 500 Example Data 169
8.2 Bayesian Formulation 170
8.2.1 The Likelihood Function 170
8.2.2 Prior Specification 170
8.2.3 The Posterior Distribution 171
8.2.4 Connection to Classical Calibration 172
8.2.5 The Prior as Regularisation: The Dupire Ill-Posedness . . . 172
8.3 BFGS as MAP Estimator and Warm Start 174
8.3.1 From Newton's Method to Quasi-Newton 174
8.3.2 BFGS Without Analytical Gradients 174
8.3.3 BFGS and the Laplace Approximation 175
8.3.4 BFGS as Warm Start for MCMC 175
8.4 Simulated Annealing and the Metropolis Algorithm 176
8.4.1 The Multi-Modal Calibration Landscape 176
8.4.2 Boltzmann Acceptance and the Metropolis Algorithm . . . 177
8.4.3 Cooling Schedule 178
8.4.4 Connection to the Bayesian Posterior 178
8.5 Metropolis-Hastings Algorithm 180
8.5.1 The M-H Kernel 180
8.5.2 Proposal Design for the Quantum Model 181
8.5.3 Burn-in and Convergence Diagnostics 182
8.5.4 Python Implementation 182
8.6 Hamiltonian Monte Carlo 186
8.6.1 From Random Walk to Dynamics 186
8.6.2 The Leapfrog Integrator 187
8.6.3 NUTS: No-U-Turn Sampler 187
8.6.4 Gradient of the Quantum Pricing Formula 188
8.7 Uncertainty Propagation to Prices and Greeks 188
8.7.1 The Posterior Predictive Distribution 188
8.7.2 Credible Intervals and Model Risk 189
8.7.3 Uncertainty in Greeks 189
8.8 Worked Example: Posterior of $\hbar_f$ from S&P 500 Smile 190
8.8.1 Data and Setup 190
8.8.2 MCMC Results 190
8.8.3 Comparison with the Point-Estimate Approach 191
8.8.4 The Posterior of $\mathcal{C}$ as a Market-State Diagnostic 192
8.9 Chapter Summary 193

IV Numerical Integration 195

9 Gaussian Quadrature for Wave Function Integrals 197
9.1 Gauss-Hermite Quadrature . . . 198
9.1.1 The General Framework . . . 198
9.1.2 Hermite Polynomials and Their Zeros . . . 199
9.1.3 The Black-Scholes Benchmark . . . 199
9.1.4 Two-Component Superposition . . . 201
9.1.5 Convergence Study . . . 201
9.1.6 The Calibrated Two-Signal State . . . 202
9.2 Gauss-Legendre for Bounded Domains . . . 204
9.2.1 Barrier Options and the Truncated Domain . . . 204
9.2.2 Legendre Polynomials and the `gauleg` Routine . . . 204
9.2.3 Gauss-Hermite versus Gauss-Legendre for Barrier Problems 205
9.3 Adaptive Quadrature . . . 206
9.3.1 The Problem of Kinks . . . 206
9.3.2 The Gauss-Kronrod Error Estimator . . . 207
9.3.3 Application to Vanilla and Exotic Payoffs . . . 208
9.3.4 Convergence Comparison . . . 209
9.4 Multi-Dimensional Quadrature . . . 209
9.4.1 The Curse of Dimensionality . . . 209
9.4.2 Tensor-Product Gauss-Hermite for the Entangled Wave Function . . . 210
9.4.3 Sparse Grids and the Smolyak Construction . . . 211
9.4.4 Break-Even Analysis . . . 212
9.5 Quadrature by Variable Transformation . . . 212
9.5.1 The Key Idea: Making the Endpoints Disappear . . . 212
9.5.2 The TANH Rule . . . 213
9.5.3 The Double Exponential (DE) Rule . . . 213
9.5.4 Infinite Ranges and the Mixed Rule . . . 214
9.5.5 The Sinc Expansion and the Wave-Theoretic Foundation . 215
9.5.6 Application: Quadrature of the Characteristic Function . . 216
9.6 Chapter Summary . . . 217

10 Oscillatory Integral Techniques 221
10.1 The Filon Method . . . 222
10.1.1 Why naive DFT fails for large ω . . . 222
10.1.2 Polynomial interpolation and the correction factor . . . 222
10.1.3 Endpoint corrections for the cubic case . . . 224
10.1.4 Practical algorithm for C_{int} . . . 224
10.1.5 Error analysis . . . 227
10.2 Levin's Method . . . 227
10.2.1 ODE reformulation . . . 227
10.2.2 Collocation implementation . . . 228
10.2.3 Convergence of Levin's method . . . 230

10.2.4 Extension to general oscillators 230
10.3 Contour Deformation and Steepest Descent 231
10.3.1 The steepest-descent principle 231
10.3.2 Saddle-point and steepest-descent path 231
10.3.3 Asymptotic expansion for large Δk_{eff} 232
10.3.4 Numerical contour deformation 232
10.4 Comparison and Method Selection 234
10.5 Euler-Maclaurin Foundation and Richardson Extrapolation 235
10.5.1 The Euler-Maclaurin summation formula 235
10.5.2 Richardson-Romberg extrapolation (NR §4.3) 235
10.5.3 Why oscillatory integrands undermine Romberg 236
10.5.4 Romberg-Filon: Richardson applied to the Filon output . . 236
10.6 Complex-Plane Path Integration of the Characteristic Function . . 237
10.6.1 Branch cuts in the Heston characteristic function 237
10.6.2 Riccati ODE integration for complex ξ 238
10.6.3 The Feynman path integral and the financial propagator . . 239
10.7 Worked Example: C_{int} for Large Δk 240
10.7.1 Setup 240
10.7.2 Reference values and convergence 240
10.7.3 Visualisation of the difficulty 241
10.7.4 Practical recommendation for the wave-mechanics pricing library 242
10.8 Chapter Summary 243

V Calibration in Practice 245

11 Equity Calibration in Practice 247
11.1 Market Data and Pre-Processing 248
11.1.1 Data Sources and Quote Conventions 248
11.1.2 Forward Price Extraction 249
11.1.3 Arbitrage-Free Conditions 249
11.1.4 Historical Volatility as the Classical Baseline 250
11.2 Loss Function Design 251
11.2.1 Weighted Implied-Volatility Residuals 251
11.2.2 Weight Choices 251
11.2.3 Tikhonov Regularisation 252
11.2.4 Maximum-Likelihood Interpretation 253
11.3 Optimisation Algorithm 253
11.3.1 The Two-Stage Strategy 253
11.3.2 Nelder-Mead Downhill Simplex 254
11.3.3 BFGS Quasi-Newton Local Refinement 255
11.3.4 Levenberg-Marquardt with Analytical Sensitivity 255
11.3.5 Convergence Criteria and Diagnostics 258
11.3.6 Confidence Regions on Calibrated Parameters 259

11.4 Case Study: SVI Surface and S&P 500 (2026) 260
11.4.1 Data and Setup 260
11.4.2 Calibrated Parameters 260
11.4.3 Local Volatility Robustness and the Quantum Density . . . 263
11.4.4 Smile Reconstruction Quality 264
11.4.5 SVI Benchmark: Industry Standard for the Smile 264
11.4.6 Real Data: SPX, 11 June 2026, and the Coherence Diagnostic 266
11.5 Stability and Robustness Tests 267
11.5.1 Sensitivity to Data Perturbation 267
11.5.2 Day-to-Day Parameter Stability 268
11.5.3 Comparison with Microstructure-Based $\hbar_f$ 268
11.5.4 Hedging Error from Miscalibrated $\hbar_f$ 269
11.6 Chapter Summary 270

12 Interest Rate Calibration in Practice 273
12.1 Market Instruments and Data 273
12.1.1 The Rate Calibration Hierarchy 273
12.1.2 Bloomberg and Reuters Identifiers 274
12.1.3 Discount Curve Bootstrapping 274
12.1.4 The Swaption Volatility Cube 277
12.2 The Quantum Short-Rate Model 277
12.2.1 Classical Background: Vasicek and CIR 277
12.2.2 Quantum Rate Model 278
12.2.3 Bond Price as Path Integral 279
12.2.4 The HJM Framework and the Quantum Forward Curve . . 279
12.2.5 Bond Options under Vasicek 282
12.2.6 Swap Rate and Level Function 282
12.3 Calibration to the Volatility Cube 283
12.3.1 Stage 1: Yield Curve Fit (Initial Condition) 283
12.3.2 Stage 2: Caplet Strip (Variance Term Structure) 283
12.3.3 SABR as Smile Benchmark 284
12.3.4 Stage 3: Swaption Surface (Correlation Structure) 285
12.3.5 Semidefinite Programming Calibration 286
12.4 Case Study: Euribor Swaption Surface 289
12.4.1 Data and Setup 289
12.4.2 Stage 1: Yield Curve 289
12.4.3 Stage 2: Caplet Strip and $\hbar_f^r$ Identification 290
12.4.4 Stage 3: Swaption Cube Fit 291
12.4.5 Principal Component Analysis of the Yield Curve 291
12.5 Negative Rates and Tunnelling Through Zero 293
12.5.1 The Zero Lower Bound as a Quantum Barrier 293
12.5.2 CIR Tunnelling Probability 293
12.5.3 ZLB Barrier Calibration: ECB 2014–2022 294
12.5.4 Tracking Error from Miscalibrated $\hbar_f^r$ 295

12.6 Chapter Summary 296

13 Credit Calibration: Single Names, Portfolio Loss, and FFT 299

13.1 Single-Name Calibration: The Hazard Rate Curve 300
13.1.1 Default intensity and survival probability 300
13.1.2 CDS mechanics and the par spread 301
13.1.3 Bootstrapping the hazard rate curve 301
13.1.4 Structural models: Merton calibration 303
13.1.5 First-passage-time model and Credit Grade 304
13.1.6 Fractional recovery conventions: RFV versus MVR 306
13.2 The Gaussian Copula and Correlated Defaults 307
13.2.1 Copula functions and Sklar's theorem 307
13.2.2 Single-factor model 308
13.2.3 Conditional independence 308
13.2.4 Structural correlation: the one-factor Gaussian model . . . 308
13.2.5 Intensity correlation 309
13.2.6 Default-time correlation via copula 310
13.2.7 Alternative copulas 311
13.2.8 Basket CDS and nth-to-default products 311
13.3 Portfolio Loss Distribution via FFT: The ASB Algorithm 313
13.3.1 Characteristic function of the portfolio loss 313
13.3.2 Loss grid discretisation 313
13.3.3 CDO tranche loss mechanics 314
13.3.4 The Andersen–Sidenius–Basu algorithm 314
13.3.5 Tranche pricing from the loss distribution 315
13.3.6 Complexity and speedup 318
13.4 Correlation Calibration 318
13.4.1 The compound correlation smile 318
13.4.2 Base correlation (McGinty et al., 2004) 320
13.4.3 Full calibration workflow 321
13.4.4 Degenerate cases 321
13.5 Extensions 322
13.5.1 Itô's formula with jumps and the credit PDE 322
13.5.2 CVA pricing via FFT 323
13.5.3 Dynamic credit models 323
13.5.4 Quantum perspective: default intensity as a wave function . 323
13.6 Chapter Summary 324

14 Cross-Asset Calibration and Consistency of $\hbar_f$ 327

14.1 The Cross-Asset $\hbar_f$ Hypothesis 327
14.1.1 The Multi-Asset Framework 327
14.1.2 The Financial Planck Constant as a Market-Wide Scalar . . 328
14.1.3 Operationalising the Hypothesis 329
14.1.4 The Multi-Asset Financial Schrödinger Equation 330
14.2 Joint Calibration Framework 331

14.2.1 Change of Numéraire as the Unifying Principle 331
14.2.2 Change of Numéraire as a Rescaling of the Wave Function . 332
14.2.3 Shared Parameters and Joint Loss Function 334
14.2.4 Information-Weighted Scheme 335
14.2.5 Algorithm 335
14.3 The Correlation Between $\hbar_f$ and the VIX 338
14.3.1 The VIX as a Market-Wide Stress Proxy 338
14.3.2 Does $\hbar_f$ Lead or Lag VIX? 338
14.3.3 Three Episodes and the Signatures to Look For 340
14.4 Cross-Asset Arbitrage Detection 341
14.4.1 The Quanto Adjustment Analogy 341
14.4.2 The Cross-Currency Wave Equation and Covariant Derivative 342
14.4.3 Quanto Volatility and Correlation 343
14.4.4 The Cross-Asset Signal 344
14.4.5 Three Dislocations and What the Signal Should Say 344
14.5 Chapter Summary 345

VI Implementation and Production 349

15 The Financial Wave Mechanics Online Platform 351
15.1 Why an Online Platform Rather Than Printed Code 351
15.2 Platform Contents 352
15.3 How to Use the Platform with This Book 354
15.4 Access, Licencing, and Updates 355
15.5 Technical Architecture 355
15.6 REST API 356
15.7 Validation and Test Suite 358
15.8 Chapter Summary 360

16 GPU Acceleration 363
16.1 GPU Architecture for Financial Computation 363
16.1.1 SIMD vs. SIMT 363
16.1.2 Memory Hierarchy and Bandwidth Bottleneck 364
16.2 CuPy and JAX for Quantum Finance 365
16.2.1 CuPy: Drop-In NumPy Replacement 365
16.2.2 JAX: JIT Compilation and Automatic Differentiation . . . 367
16.3 GPU-Accelerated Monte Carlo 369
16.3.1 Random Number Generation on GPU 369
16.3.2 Quantum Reweighting on GPU 370
16.3.3 Multi-Asset Monte Carlo 371
16.4 GPU-Accelerated Crank-Nicolson 372
16.4.1 The Parallelisation Challenge 372
16.4.2 Strategy 1: Batch Solve Across Options 372
16.4.3 Strategy 2: Parallel Cyclic Reduction 373

16.5 Benchmark Results . . . 375
16.5.1 Hardware and Software Environment . . . 375
16.5.2 Speedup Table . . . 375
16.5.3 Cloud Cost Estimates . . . 376
16.6 Chapter Summary . . . 376

17 Backtesting Framework **379**
17.1 Backtesting Infrastructure . . . 380
17.1.1 What the Framework Must Store . . . 380
17.1.2 Data Pipeline . . . 380
17.1.3 Daily Calibration Run . . . 381
17.1.4 Out-of-Sample Protocol . . . 382
17.1.5 Database Schema . . . 382
17.2 Option Pricing Performance . . . 383
17.2.1 Metric: Out-of-Sample Implied Volatility RMSE . . . 383
17.2.2 Benchmark Models . . . 383
17.2.3 What Has Been Measured . . . 383
17.2.4 What the Regime Study Must Report . . . 384
17.2.5 The Rates Market: a Structural Argument . . . 385
17.3 Delta Hedging Performance . . . 385
17.3.1 Experimental Design . . . 385
17.3.2 The Quantum Delta . . . 386
17.3.3 Illustration: How Large Is the Delta Gap? . . . 387
17.4 VaR Backtesting . . . 389
17.4.1 Setup . . . 389
17.4.2 Statistical Tests . . . 389
17.4.3 The Basel Traffic Light . . . 390
17.4.4 Illustration: Can the Tests Detect a Bad Model? . . . 390
17.5 Coherence as the Market State Gauge . . . 394
17.5.1 Which Calibrated Quantity Is the Gauge? . . . 394
17.5.2 The Lead Hypothesis and Its Test . . . 395
17.5.3 Coherence as a Crisis Predictor: the ROC Protocol . . . 395
17.6 Chapter Summary . . . 396

18 Real-Time Systems and Production Deployment **399**
18.1 Latency Requirements by Use Case . . . 399
18.1.1 The Latency Budget . . . 399
18.1.2 Market Making: Sub-Millisecond Pricing . . . 399
18.1.3 Intraday Risk: Warm-Start Recalibration . . . 400
18.2 Sub-Millisecond Calibration . . . 401
18.2.1 Architecture . . . 401
18.2.2 Perturbative Update . . . 401
18.2.3 Benchmark: From 50ms to 0.8ms . . . 402
18.3 Real-Time $\hbar_f$ and Coherence Monitor . . . 405
18.3.1 The $\hbar_f$ and Coherence Time Series . . . 405

18.3.2 Coherence, Contrast, and the Coherence Time 405
18.3.3 Alert Thresholds 406
18.3.4 Cross-Asset Dashboard 406
18.4 Integration with Risk Systems 408
18.4.1 FRTB Internal Model Approach 408
18.4.2 Quantum CVA 409
18.4.3 Portfolio-Level Quantum Greeks 409
18.5 Model Governance and Validation 410
18.5.1 Regulatory Framework 410
18.5.2 Validation Test Battery 410
18.5.3 Model Risk Reserve and the Planck Greek 412
18.5.4 Documentation and Governance Process 412
18.6 Chapter Summary and Closing Remarks 413

Preface to Volume II

Volume I of this work did two things. It developed the theoretical framework of financial wave mechanics: the financial Schrödinger equation and its wave phenomena. It then applied that framework to every major asset class and to every major risk management problem. And it took the framework to the data. Volume I prices a live CBOE SPX option chain and fits three decades of VIX history. It reads the June 2026 US Treasury curve as a wave function. It measures quantum value at risk against fifty years of S&P 500 returns. Volume II asks the complementary question: *how do you compute all of this at production scale?*

This is not a minor question. A closed-form quantum option price is beautiful, but a beautiful formula is not yet a pricing library. Between the two lie the choice of quadrature, the damping parameter, the grid, the optimiser and the stopping rule. The practitioner who has read Volume I asks a fair question. "Very well, but how do I price this barrier option with a quantum correction on Monday morning?" It deserves a complete and honest answer.

This volume is that answer. It is organised as a systematic treatment of the numerical methods that suit the wave-mechanical setting best: eighteen chapters in six parts. Part I begins with the most analytically tractable: Fourier methods and perturbation theory. Part II turns to grid-based solvers, from finite differences to finite elements and spectral methods. Part III covers Monte Carlo, including quasi-random sequences and Markov chain calibration. Part IV treats numerical integration itself: Gaussian quadrature, and the oscillatory integrals that interference makes unavoidable. Part V works through calibration asset class by asset class — equity, rates, credit — and then across them, where one $\hbar_f$ is shared. Part VI closes with the production tools — the online platform, GPU acceleration, backtesting and real-time systems.

The title *Numerical Recipes for Financial Wave Mechanics* is chosen deliberately. Like the classic numerical analysis reference it invokes, this volume aims to be a book that practitioners keep open alongside their code: not a textbook to be read from cover to cover, but a handbook to be consulted when a specific computational problem arises. Each chapter presents the method, its complexity, its accuracy guarantees, its known failure modes, and its implementation recipe — a self-contained guide to applying the technique to the quantum finance problems of Volume I. Each recipe names the Volume I result it implements, gives the algorithm step by step with its stopping rule, and ends with a complete Python listing, a worked example on the calibrated state, the pitfalls, a measured cost and a verdict on when to use it. Every listing was run for this edition, and each timing names the machine that produced it.

A word on data. The corpus behind the two volumes holds one live option snapshot, the CBOE S&P 500 chain of 11 June 2026, and three public series: the VIX, the US Treasury curve and fifty years of S&P 500 closes. Where a chapter uses licensed data it says so, and where a study has not yet been run it prints the template rather than a number.

The online platform described in Chapter 15 (`https://quantum.abdelkader.site`) provides interactive implementations of every method in this volume: online calculators, documented notebooks, and calibrated datasets. The platform is maintained and updated continuously, so that the code never goes stale with the book.

I have written this volume for the reader who has to make the framework work. That means the quantitative analyst building the pricing library, the risk manager who must validate it, and the graduate student who wants to reproduce every number of Volume I. Some familiarity with numerical analysis helps, but I have assumed none. Each method is developed from its first principle, with the worked example and the failure modes that a working implementation needs. Volume I is the natural companion, although a reader who only needs one recipe can start with the chapter that contains it.

My thanks go again to the students, the traders, the risk managers and the colleagues whose questions shaped Volume I. This volume grew out of the same conversations. I owe a particular debt to the practitioners who kept asking, of every result, how the number had actually been obtained. Their insistence set the agenda for these chapters. Any errors that remain are mine alone.

Abdelkader Bousabaa
Paris, August 2026

A beautiful formula is not yet a pricing library. This is the code behind the theory.

Volume I derived the financial Schrödinger equation and took it to the data. This volume answers the next question: how do you compute all of it, correctly and fast enough to quote on Monday morning? Eighteen chapters, each with a step-by-step recipe, a complete Python listing that runs, a worked example on the calibrated S&P 500 state, the pitfalls that actually bite, a measured cost and a verdict on when to use the method.

Volume II covers the numerical methods that suit wave mechanics best, in six parts: Fourier transforms and the quantum characteristic function; perturbation theory and WKB; finite differences, finite elements and spectral solvers that stay unitary; Monte Carlo, quasi-random sequences and Bayesian calibration by MCMC; Gaussian quadrature and the oscillatory integrals that interference makes unavoidable; calibration asset class by asset class and then across them, where one financial Planck constant is shared; and the production stack — online platform, GPU acceleration, backtesting protocol and real-time systems that recalibrate in under a millisecond. Every listing was run for this edition, every timing names the machine that produced it, and where a study has not been run the book prints the template rather than a number.

Abdelkader Bousabaa is a quantitative finance practitioner and lecturer based in Paris. Trained at Crédit Lyonnais, he has taught quantitative finance for more than twenty-five years and worked as a consultant for banks, asset managers and insurers on derivatives pricing, model validation and risk management. He founded and developed an asset management company in London, and has built and validated quantitative models in production.

Seven Lions Editions