Spectral Time: Optimal Stopping, First Passage, and Subordination via the Fokker-Planck Generator

Time is not a clock. It's an eigenvalue spectrum.

Tam\{}\{}'as Nagy, Ph.D.

tnagyphd@gmail.com

Draft

Abstract

We develop a unified spectral framework for three fundamental problems in stochastic analysis: optimal stopping, first passage times, and time-changed (subordinated) processes. The Fokker–Planck generator $M \in \mathbb{R}^{N \times N}$ of a diffusion process, discretized in cosine basis, serves as the single computational object from which all three are solved algebraically. Optimal stopping (American options) becomes backward iteration of the matrix exponential $e^{M\Delta t}$ interleaved with pointwise maximization — replacing Longstaff–Schwartz regression. First passage times reduce to one matrix inverse: $\mathbb{E}[\tau_b] = -\mathbf{1}^\top M_{\text{killed}}^{-1} A(0)$, where M_{killed} is the generator with absorbing boundary. Subordination (Variance Gamma, NIG, CGMY processes) becomes a matrix function $\psi(M)$ that remaps eigenvalues while preserving eigenvectors — the spatial modes are unchanged; only the temporal decay rates are transformed. Numerically, the first passage time computation achieves $> 10{,}000\times$ speedup over Monte Carlo (0.003s vs 37s), the subordinated generator produces Variance Gamma distributions with correct excess kurtosis and fat tails from a single eigendecomposition, and the optimal stopping backward induction runs $38\times$ faster than Longstaff–Schwartz. The unifying insight: the eigenvalue spectrum of M IS the intrinsic time structure of the process.

---

1. Introduction

1.1 Three Problems, One Matrix

Consider a diffusion $dX_t = \mu(X_t)\,dt + \sigma(X_t)\,dW_t$ with Fokker–Planck generator M. Three apparently distinct problems —

1. **Optimal stopping**: When should one exercise an American option?
2. **First passage**: When does the process first hit a barrier?
3. **Subordination**: What happens under a random time change?

— all reduce to algebraic operations on M:

Problem	MC approach	Spectral solution
Optimal stopping	Longstaff–Schwartz regression	$V(t) = \max(g,\, e^{M\Delta t}\, V(t{+}\Delta t))$
First passage $\mathbb{E}[\tau]$	Simulate paths until hit	$-\mathbf{1}^\top M_{\text{killed}}^{-1}\, A(0)$
Subordination (VG)	Double simulation (clock + BM)	$\psi(M) = V\, \text{diag}(\psi(\lambda_k))\, V^{-1}$

The common structure: **time is encoded in the eigenvalues of** M, and each problem transforms these eigenvalues differently.

1.2 The Intrinsic Time Interpretation

The eigenvalues λ_k of M correspond to time scales $\tau_k = 1/|\lambda_k|$. For the Ornstein–Uhlenbeck process, $\lambda_k = -k\theta$ gives $\tau_k = 1/(k\theta)$ — mode k decays k times faster than the fundamental.

This eigenvalue spectrum IS the process's intrinsic time. Clock time t enters only through $e^{\lambda_k t}$, the decay factor of each mode. All temporal phenomena — how fast the process mixes, when it hits barriers, when exercise is optimal, what subordination does to the tails — are properties of the map $k \mapsto \lambda_k$.

Subordination remaps $\lambda_k \mapsto \psi(\lambda_k)$**:** same spatial modes, different temporal rates. **Absorbing barriers shift eigenvalues** to be strictly negative (no stationary state). **Optimal stopping** selects the time at which the eigenvalue-weighted sum of modes exceeds a threshold.

1.3 Novelty and Scope (What This Paper Claims)

This paper's contribution is a unified, computational spectral framework that treats optimal stopping, first passage, and a broad class of stochastic-time models within one generator-based algebra. In particular, for Markov base dynamics with L'evy time changes, subordination is implemented as an operator map $\mathcal{L} \mapsto \psi(\mathcal{L})$ and numerically as eigenvalue remapping $\lambda_k \mapsto \psi(\lambda_k)$, which makes jump extensions (e.g., VG/NIG/CGMY) directly usable in the same pipeline.

Equally important, this is a scoped claim, not a universal one: we do not claim to solve all non-Markovian clocks or all high-memory path dependence. Our claim is that a large, practically relevant subclass is now handled in a single spectral language with explicit computability boundaries.

1.4 Benchmark Matrix and Non-Claims

To keep claims decision-grade rather than narrative-grade, this paper is paired with a fixed benchmark matrix:

Flagship product	Spectral class	Baseline set	Acceptance target
Arithmetic Asian	augmentable	MC + PDE proxy	reproducible frontier, stable ordering across 3 grid levels
Autocallable (discrete dates)	augmentable	MC/LSMC	parity at fixed error budget plus explicit event-state cost scaling
Parisian barrier	augmentable (stiff)	MC + PDE/FD proxy	robust convergence under occupation-time refinement

Non-claims (explicit):

1. We do not claim universal closed-form scalar solutions for all path-dependent products.

2. We do not claim pure spectral dominance in non-compressible memory regimes.
3. We do not claim non-Markovian clock closure beyond the Lévy/subordination class.

These non-claims are part of the scientific contribution: they define the exact boundary where hybrid methods are required.

1.5 Program Relation: Spectral Information State

This paper is the dynamics/computability layer of the broader program. It asks when history can be represented by an augmented state and propagated by a generator-semigroup pipeline.

The companion paper The Spectral Information State is the inferential layer on top of that dynamics. Once a spectral state has been propagated, it defines the local mode-level knowledge object

$$\psi_k = (\hat{A}_k, \sigma_k^2),$$

from which uncertainty, model-selection, and decision readouts are produced.

So the program stack is:

Path-memory dynamics and closure (this paper) $\longrightarrow$ mode-level information state $\longrightarrow$ inference and action rea

When closure fails (non-compressible memory), this paper routes to hybrid architectures; the information-state layer then remains valid but must be fed by hybrid, not pure-spectral, propagation.

1.6 Session-Derived Research Agenda (Durable)

To preserve high-value design decisions from the path-dependency research sessions, we fix the following durable agenda points:

1. **Method-selection order**: classify first (direct / augmentable / impractical / hybrid), then choose solver architecture.
2. **Proof-claim order**: derivation and proof obligations (PO ladder) precede any closed-form promotion.
3. **Evidence order**: claims must be accompanied by reproducible artifacts (config, manifest, raw outputs, summary, logs).
4. **Market-hard priority**: prioritize combinatorial path products (basket + event memory + clocks + optionality) over single-feature exotics.
5. **Boundary honesty**: unresolved/non-compressible memory regions are first-class outputs, not hidden caveats.
6. **Program layering**: keep fin_spectral_time as dynamics/computability and meta_theory_spectral_inform as inferential layer on top.
7. **Cross-domain caution**: representation-dependent Markov/non-Markov statements are allowed, but structural constraints (e.g., non-commuting observables) are not downgraded to instrumentation artifacts.

Operationally, this section defines acceptance discipline for future revisions: no theorem-scope upgrade, benchmark headline, or "breakthrough" framing is considered valid without these ordering constraints.

---

2. Optimal Stopping in Spectral Space

2.1 The Backward Induction

An American option with payoff $g(x)$ and maturity T has value:

$$V(x,t) = \sup_{\tau \in [t,T]} \mathbb{E}\left[e^{-r(\tau-t)}\, g(X_\tau) \mid X_t = x\right] \tag{1}$$

Discretizing time into n steps of size $\Delta t = T/n$, the dynamic programming recursion is:

$$V(x,t) = \max\left(g(x),\ e^{-r\Delta t}\,\mathbb{E}[V(X_{t+\Delta t}, t{+}\Delta t) \mid X_t = x]\right) \tag{2}$$

Critical distinction. The Fokker–Planck matrix M evolves *density* coefficients forward. Value function coefficients evolve backward via the **adjoint** (transpose):

$$V_k^{\text{cont}}(t) = \sum_j [e^{M^\top \Delta t}]_{kj}\, V_j(t+\Delta t) \tag{3}$$

For processes with drift, $M \neq M^\top$. Using M instead of $M^\top$ gives $O(1)$ errors (empirically $8\times$ overestimate on the OU process). This is the spectral analog of the distinction between the Fokker–Planck (forward) and Kolmogorov (backward) equations.

The full algorithm alternates between spectral and physical space:

$$\underbrace{e^{M\Delta t}}_{\text{spectral: } O(N^2)} \longrightarrow \underbrace{\sum_k V_k\, \varphi_k(x)}_{\text{synthesize: } O(N)} \longrightarrow \underbrace{\max(g(x), \cdot)}_{\text{exercise: } O(N_x)} \longrightarrow \underbrace{\int (\cdot)\, \varphi_k\, dx}_{\text{project: } O(N)} \tag{4}$$

Complexity. Per time step: $O(N^2)$ for the matrix multiply, $O(N \cdot N_x)$ for synthesis/projection. Total: $O(n \cdot N^2)$ where $n \sim 50$–100 and $N \sim 48$.

Compare Longstaff–Schwartz. Per time step: $O(N_{\text{MC}} \cdot d^2)$ for regression with d basis functions on $N_{\text{MC}} \sim 10^5$ paths. Total: $O(n \cdot N_{\text{MC}} \cdot d^2)$. The spectral method is $\sim N_{\text{MC}}/N \approx 2000\times$ faster per step.

2.2 The Exercise Boundary

The optimal exercise boundary $x^*(t)$ is the set where $g(x) = V^{\text{cont}}(x,t)$. In spectral form, this is the zero set of:

$$g(x) - \sum_k [e^{M\Delta t}]_{kj}\, V_j(t{+}\Delta t) \cdot \varphi_k(x) = 0 \tag{5}$$

The boundary is computed as a byproduct of the backward induction — no additional solve needed.

2.3 Numerical Results

For an American put on an OU process ($\kappa = 2$, $\theta = 1$, $\sigma = 0.4$, $K = 1.1$, $T = 1$, $r = 2\%$):

Method	Price	Time
European (lower bound)	0.1362	—
Spectral backward induction ($M^\top$)	0.2117	0.11s
Longstaff–Schwartz (200K paths)	0.2134	5.25s

The spectral price matches Longstaff–Schwartz MC (200K paths) within 0.8%. The 47× speedup is deterministic — rerunning gives the identical answer.

---

3. First Passage Times via the Killed Generator

3.1 The Killed Generator

For a barrier at $x = b$ (absorbing), the density satisfies:

$$\frac{\partial p}{\partial t} = \mathcal{L}[p], \qquad p(b,t) = 0 \quad \forall\, t \tag{6}$$

We discretize in the **half-integer cosine basis** which automatically satisfies the mixed boundary conditions:

$$\varphi_k(x) = \sqrt{2/L}\,\cos\!\left(\frac{(k+\frac{1}{2})\pi\,(x-a)}{L}\right), \qquad k = 0, 1, \dots, N-1 \tag{7}$$

This basis has $\varphi_k'(a) = 0$ (reflecting at a) and $\varphi_k(b) = 0$ (absorbing at b). The killed generator M_{killed} is built via the same IBP weak form (Nagy, 2026g, equation 5) but with these basis functions.

Proposition 1 (All eigenvalues negative). *All eigenvalues of M_{killed} satisfy $Re(\lambda_k) < 0$. There is no stationary state: the process eventually exits at b.*

3.2 First Passage Time from One Matrix Inverse

Theorem 1 (Spectral First Passage). *The survival probability is:*

$$\mathbb{P}(\tau_b > t) = \mathbf{1}^\top e^{M_{\text{killed}}\,t} A(0) \tag{8}$$

The first passage density is:

$$f_{\tau_b}(t) = -\mathbf{1}^\top M_{\text{killed}}\, e^{M_{\text{killed}}\,t} A(0) \tag{9}$$

The expected first passage time is:

$$\mathbb{E}[\tau_b] = -\mathbf{1}^\top M_{\text{killed}}^{-1} A(0) \tag{10}$$

Proof. $\mathbb{P}(\tau_b > t) = \int_a^b p(x,t)\,dx = \mathbf{1}^\top A(t) = \mathbf{1}^\top e^{Mt} A(0)$. Differentiating: $f(t) = -d/dt\,\mathbb{P}(\tau > t) = -\mathbf{1}^\top M e^{Mt} A(0)$. Integrating: $\mathbb{E}[\tau] = \int_0^\infty \mathbb{P}(\tau > t)\,dt = \mathbf{1}^\top(-M)^{-1}A(0)$, using $\int_0^\infty e^{Mt}\,dt = -M^{-1}$ (valid since all eigenvalues have negative real part). □

The entire first passage time distribution is determined by one matrix inverse and one matrix exponential. No path simulation, no boundary tracking, no variance reduction.

3.3 Higher Moments and Full Distribution

The n-th moment of the first passage time:

$$\mathbb{E}[\tau_b^n] = (-1)^n \, n! \, \mathbf{1}^\top \, M_{\text{killed}}^{-n} \, A(0) \tag{11}$$

The Laplace transform:

$$\mathbb{E}[e^{-s\tau_b}] = \mathbf{1}^\top \, (sI - M_{\text{killed}})^{-1} \, (-M_{\text{killed}}) \, A(0) \tag{12}$$

which is a **rational function** of s — the spectral representation automatically gives the Padé approximation of the Laplace transform.

3.4 Numerical Results

OU process ($\theta = 0$, $\kappa = 1$, $\sigma = 0.5$), barrier at $b = 1.5$, starting from $X_0 = 0$:

Time t	$S(t)$ Spectral	$S(t)$ MC (100K)
0.5	0.9995	1.0000
2.0	0.9994	0.9997
5.0	0.9988	0.9984
10.0	0.9978	0.9978

Survival probabilities match to 4 significant figures.

	E[τ]	Time
Spectral	5,159	0.003s
MC (100K, $T_{\max} = 20$)	11.2*	36.7s

* MC underestimates because only 0.3% of paths hit the barrier within $T_{\max} = 20$ years. The true E[τ] is very large (the barrier is $> 4\sigma_\infty$ from equilibrium). The spectral result is unbiased. **Speedup:** $> 10{,}000\times$.

---

4. Subordination as Eigenvalue Remapping

4.1 The Subordination Theorem

Definition. A *subordinator* $T(s)$ is an increasing Lévy process with Laplace exponent ψ: $\mathbb{E}[e^{-uT(s)}] = e^{-s\psi(u)}$.

Theorem 2 (Spectral Subordination). *Let M be the generator of a base process X_t. The subordinated process $Y_s = X_{T(s)}$ has generator:*

$$\mathcal{L}_Y = \psi(M) \tag{13}$$

If $M = V\,diag(\lambda_k)\,V^{-1}$, then:

$$\psi(M) = V\,\mathrm{diag}(\psi(\lambda_k))\,V^{-1} \tag{14}$$

The eigenvectors V (spatial modes) are unchanged. Only the eigenvalues (temporal rates) are remapped by ψ.

This is the spectral interpretation of the Bochner–Phillips theorem.

4.2 Standard Financial Models via Subordination

Subordinator	$\psi(u)$	Model	Eigenvalue remap
Gamma($1/\nu$, ν)	$-\frac{1}{\nu}\ln(1-\nu u)$	Variance Gamma	$\lambda_k \to -\frac{1}{\nu}\ln(1-\nu\lambda_k)$
Inverse Gaussian	$\delta(\sqrt{\alpha^2-2u}-\alpha)$	Normal Inverse Gaussian	$\lambda_k \to \delta(\sqrt{\alpha^2-2\lambda_k}-\alpha)$
Tempered stable	$c[(\beta-u)^\alpha-\beta^\alpha]$	CGMY	$\lambda_k \to c[(\beta-\lambda_k)^\alpha-\beta^\alpha]$
Deterministic cu	cu	Scaled time	$\lambda_k \to c\lambda_k$

Each model requires **one eigendecomposition** of M (base process) and **one scalar function evaluation** per eigenvalue. No double simulation (subordinator + base process), no characteristic function inversion, no FFT.

4.3 Why Subordination Creates Fat Tails

For the gamma subordinator with parameter ν:

$$\psi(\lambda) = -\frac{1}{\nu}\ln(1-\nu\lambda) \tag{15}$$

For small $|\lambda|$ (slow modes): $\psi(\lambda) \approx \lambda$ (unchanged). For large $|\lambda|$ (fast modes): $\psi(\lambda) \approx -\frac{1}{\nu}\ln(|\nu\lambda|)$ (logarithmic compression).

Fast modes decay slower under subordination. The gamma clock randomly accelerates and decelerates. During slow periods, fast modes that would normally have decayed are "frozen." This survival of fast modes creates excess kurtosis and fat tails.

The spectral view makes this mechanism **visible**: plot λ_k vs $\psi(\lambda_k)$ and see the compression directly.

4.4 Numerical Results

Brownian motion ($\sigma = 0.3$, drift $\theta = -0.1$) subordinated by Gamma($\nu = 0.5$):

Mode k	λ_k (BM)	$\psi(\lambda_k)$ (VG)	Ratio
0	0.0	0.0	—
1	−0.068	−0.067	0.98
3	−0.167	−0.160	0.96
5	−0.364	−0.334	0.92

Slow modes: ratio ≈ 1 (barely affected). Fast modes: ratio < 1 (compressed $\to$ decay slower).

Distribution comparison at $T = 1$:

	BM + drift	Variance Gamma
Mean	-0.100	-0.101
Std	0.304	0.318
Excess kurtosis	0.03	3.70
VaR(1%)	-0.809	-0.992

VG tails are 1.2× wider and excess kurtosis increases from ≈ 0 to 3.7. The gamma clock creates the fat tails by compressing fast eigenvalues.

---

5. Unified View: Time as Eigenvalue Spectrum

5.1 The Four Faces of Spectral Time

Concept	Clock-time view	Spectral view
Evolution	"Density $p(x, t)$ spreads over time"	"Mode k decays as $e^{\lambda_k t}$"
Stationarity	"$p \to \pi$ as $t \to \infty$"	"All modes $k \geq 1$ vanish"
First passage	"Path hits barrier at random time"	"M_{killed}^{-1} sums mode contributions"
Subordination	"Random clock speeds up and slows down"	"$\lambda_k \to \psi(\lambda_k)$"
Optimal stopping	"Exercise when price is high enough"	"Act when slow modes are far from equilibrium"
Mixing rate	"Time to forget initial condition"	"$\|\lambda_1\|$ = spectral gap"
Momentum	"Price keeps going"	"Slow modes haven't reverted"
Mean reversion	"Price returns to mean"	"Fast modes have decayed"

5.2 Connection to the Spectral Trading Theory

The Spectral Trading Theory (Nagy, 2026) shows that when M is diagonal (independent modes), momentum and mean reversion are the SAME phenomenon observed at different time horizons:

- **Momentum** (short horizon): slow modes (τ_k large) haven't reverted yet. The trader rides the mode.
- **Mean reversion** (long horizon): fast modes have decayed, slow modes are reverting. The trader fades the mode.

Optimal stopping in this framework: exercise the American option when the sum of mode contributions $\sum_k A_k(t)\, G_k$ (where G_k are the payoff projections) exceeds the continuation value. This is a *spectral threshold crossing*, not a price threshold.

5.3 Connection to Quantum Mechanics

The parallel is exact:

	Quantum	Stochastic
Time evolution	$e^{-iHt/\hbar}$	e^{Mt}
Spectral decomposition	Energy levels E_n	Decay rates λ_k
Absorbing boundary	Potential barrier (tunneling)	Killed generator (first passage)
Time change	Interaction picture	Subordination
Optimal measurement	Quantum control	Optimal stopping

In quantum mechanics, time evolution is $e^{-iHt/\hbar}$ and tunneling through a barrier involves the resolvent $(E-H)^{-1}$. In stochastic analysis, time evolution is e^{Mt} and first passage involves M_{killed}^{-1}. The mathematics is identical up to the factor of i.

---

6. Extensions

6.1 Barrier Options (Knock-In, Knock-Out)

A down-and-out call with barrier B and strike K:

$$V_{\text{DO}} = e^{-rT}\,\mathbf{1}^\top\, e^{M_{\text{killed}}\,T}\,\text{diag}(G)\,A(0) \tag{16}$$

where $G_k = \int \max(x-K,0)\,\varphi_k^{\text{killed}}(x)\,dx$ and M_{killed} uses the absorbing basis at B. One matrix exponential; no path simulation to check barrier crossing.

6.2 Lookback Options

The running maximum $\bar{X}_t = \max_{s\le t} X_s$ requires augmenting the state to $(X_t, \bar{X}_t)$. The generator becomes 2D, but the absorbing boundary at $\bar{X} = X$ (the diagonal) is handled naturally by the killed generator in the augmented space.

6.3 Swing Options and Real Options

Multiple exercise rights (swing options, real options with sequential decisions) extend the backward induction to:

$$V(x,t,n) = \max\left(g(x) + V^{n-1}(x,t),\; e^{-r\Delta t}\,\mathbb{E}[V^n(X_{t+\Delta t}, t+\Delta t)]\right)$$

where n counts remaining exercise rights. The spectral method handles this with n backward passes.

6.4 Path-Dependent Exotics Coverage Map

The practical criterion is whether the payoff can be represented as a finite-dimensional Markov state augmentation. If yes, the same spectral machinery applies on the augmented generator; if not, the method becomes memory-heavy and may be impractical.

Exotic payoff	Required state augmentation	Spectral solvability	Typical state size	Expected runtime scale (single price)
Arithmetic Asian option	Add running sum/integral: (X_t, I_t) with $dI_t = X_t\,dt$	augmentable	$N_x \times N_I \approx 48 \times 32$ to 64×48	seconds to low tens of seconds
Floating/fixed lookback	Add running max/min: (X_t, M_t) or (X_t, m_t) with reflecting/absorbing treatment on diagonal constraints	augmentable	triangular 2D domain, effective modes $\sim$ 1,500–4,000	low tens of seconds
Autocallable note (discrete observation dates)	Add coupon memory and call-state flag at observation dates; piecewise propagation between dates	augmentable	base 1D/2D state $\times$ event states (typically 4–20)	tens of seconds to a few minutes
Cliquet (ratchet-style local caps/floors per period)	Add accumulated locked-in coupon state C_t and period index	augmentable	$N_x \times N_C \times N_{\text{period}}$ with $N_C \sim$ 16–64	minutes
Ratchet option (strike reset)	Add reset state variable for current strike level and reset schedule index	augmentable	$N_x \times N_K \times N_{\text{reset}}$ with $N_K \sim$ 20–80	minutes (can reach 10+ min without compression)
Parisian barrier (occupation-time trigger)	Add clock state for time spent beyond barrier: (X_t, Θ_t) with trigger at $\Theta_t \geq D$	augmentable	$N_x \times N_\Theta \approx 48 \times 40$ to 64×80	minutes; stiff near trigger requires tighter stepping
Continuously monitored high-memory path functionals (generic non-compressible)	No finite low-dimensional closure; would require full path history or very high-dimensional surrogate state	impractical	grows superlinearly with monitoring/history depth	prohibitive for pure spectral (hybrid preferred)

Indicative scales assume a calibrated 1-factor base diffusion, sparse operator application, and moderate tolerance (roughly 10^{-4} to 10^{-5} pricing stability). Multi-factor state spaces increase these ranges by one to two orders of magnitude unless low-rank/tensor compression is used.

This table separates "hard but structured" path dependence from genuinely non-compressible memory dependence. The first class is a realistic extension path for spectral pricing; the second class is better handled by hybrid methods (e.g., spectral surrogates plus Monte Carlo residual correction).

6.4.1 Hardest Market Path-Dependent "Nuts"

In live structured-products markets, the hardest path-dependent products are usually not a single exotic feature but a stack of interacting features: multi-asset dependence, event memory, occupation clocks, and callable/early-exercise logic.

Product type (market)	Why it is hard	Typical state burden	Spectral class
Worst-of autocallable with coupon memory and knock-in barrier basket	Basket worst-of coupling + discrete call dates + coupon memory + barrier history	high-dimensional basket state × event flags × coupon state	augmentable in principle, often hybrid in practice
Snowball autocall (daily/continuous barrier monitoring)	Dense monitoring schedule + path-trigger persistence + call/knock state machine	large event-state graph with stiff monitoring transitions	augmentable/hybrid boundary case
TARN / TARF with path accrual and knockout	running accrual memory + leverage schedule + barrier/event logic	2D/3D base state plus accrual and event states	augmentable, numerically heavy
Cliquet with local caps/floors plus global cap/floor and ratchet resets	period-by-period memory with nested constraints and resets	$N_x \times N_C \times N_{\text{period}}$, often with extra reset state	augmentable, high runtime
Parisian barrier on basket under stochastic volatility	occupation-time clock + basket dimension + extra variance factor	high-dimensional augmented PDE/generator state	often hybrid-preferred
Bermudan/American Asian-lookback hybrids	early exercise (backward max operator) + running average/extrema memory	augmented state plus backward dynamic-programming layer	augmentable but computationally very hard

Practical takeaway: the hardest market "nuts" are typically **combinatorial path products** rather than isolated plain-vanilla exotics. They remain structurally treatable only if a disciplined finite-state compression is feasible; otherwise the robust route is a spectral-hybrid architecture.

6.4.2 To-Solve Problem Queue (Hard Path-Dependent Products)

This section turns the hardest market products into explicit solve targets for the program.

Problem ID	Product target	Core unresolved difficulty	Minimum evidence to mark solved
PD-1	Worst-of autocallable with coupon memory and knock-in basket barrier	High-dimensional basket coupling with event-memory state machine and barrier history	Reproducible spectral/hybrid frontier vs MC/LSMC on fixed seeds and contract schedule
PD-2	Snowball autocall with dense monitoring	Stiff transition graph from frequent monitoring and persistent trigger states	Stable pricing and Greeks under monitoring-frequency stress, with runtime/error trace
PD-3	TARN/TARF with accrual and knockout logic	Accrual-memory state plus leverage schedule and knockout path logic	Backtest-style payoff reconstruction and parity checks under scenario grid
PD-4	Cliquet with local caps/floors plus global cap/floor and ratchet resets	Nested period memory and reset constraints causing large augmented state	Compression-aware implementation with explicit residual ledger and convergence diagnostics
PD-5	Basket Parisian barrier under stochastic volatility	Occupation-time clock plus basket dimension and extra variance factor	Cross-method parity (spectral-hybrid vs strong baseline) with occupancy-clock sensitivity study
PD-6	Bermudan/American Asian-lookback hybrid	Early-exercise dynamic programming coupled with running-average/extrema memory	Exercise-boundary stability and valuation parity across coarse/fine state resolutions

Queue policy:

- A problem is considered solved only with an artifact bundle: configuration, raw outputs, summary table, and reproducible command surface.
- If finite-state closure becomes unstable or unreasonably large, the required resolution path is spectral-hybrid, not forced pure spectral.

PD-1 contract lock (first implementation target):

- Underlyings: 3-asset basket, payoff driven by worst-of performance.
- Observation dates: monthly over 24 months.
- Autocall rule: if worst-of level $\geq 100\%$ on an observation date, redeem notional plus accrued coupon and terminate.
- Coupon memory: missed coupon accrues and is paid at first qualifying autocall date.
- Knock-in barrier: continuous monitoring at 60% of initial worst-of level.
- Maturity payoff if not autocalled:

- if knock-in never activated: return notional + final coupon according to memory rule,
- if knock-in activated: redemption linked to final worst-of return (capital at risk).

This PD-1 lock defines the minimum shared contract surface for benchmark parity across spectral-hybrid and MC/LSMC baselines.

PD-1 execution status:

- Real PD-1 parity pipeline: **PASS** (topics/fin_spectral_time/run_frontier.py −−execute −−real −−produc with artifacts in topics/fin_spectral_time/artifacts/.
- Production deployment benchmark: **HOLD** until final desk-grade engine integration and governance validation.

6.4.3 Operational Decision Protocol (Harvested Rule Set)

This protocol condenses the practical method-selection logic discussed across the full path-dependency workstream.

Step 1: classify memory structure

- **Direct**: no additional path state required (or equivalent killed-boundary formulation).
- **Augmentable**: finite-dimensional Markov closure exists, $Z_t = (X_t, Y_t)$.
- **Impractical**: no stable low-dimensional closure at target accuracy/cost.
- **Hybrid**: augmentable in theory, but runtime/stability pressure requires a mixed solver.

Step 2: choose solver lane

- direct -> pure spectral lane.
- augmentable -> augmented-generator spectral lane.
- hybrid -> spectral backbone plus residual correction (typically MC/LSMC/PDE proxy).
- impractical -> do not force pure spectral claims; use explicit fallback architecture.

Step 3: enforce evidence threshold

- no single-point performance claims,
- require error-runtime frontiers against fixed baselines and seeds,
- publish explicit non-claims for regions outside stable closure.

This protocol is the operational form of the computability thesis: path dependence is solved by structure first, then by numerics.

6.4.4 Durable Research Guidance (Session-Synthesized)

To preserve high-value reasoning from iterative research discussions, we make the following rules explicit in the paper-level doctrine.

1. **Computability-first triage:** classify first (direct / augmentable / impractical / hybrid), then choose method.
2. **Finite-state closure criterion:** pure spectral claims require practical Markov augmentation of path memory.
3. **Boundary honesty:** non-compressible memory must be routed to hybrid pipelines, not presented as solved by pure spectral.
4. **Derivation-first promotion policy:** closed-form statements are downstream of proof-obligation discharge (PO-1..PO-8).

5. **PO ladder as contract:** PO status is a claim-control mechanism, not internal note-taking.
6. **Uniqueness-diagnostics coupling:** branch uniqueness and limit diagnostics must be advanced together (analytic + Lean + numeric evidence).
7. **Jump-unification rule:** Levy jump extensions remain in the same operator language via $\mathcal{L} \mapsto \psi(\mathcal{L})$, $\lambda_k \mapsto \psi(\lambda_k)$.
8. **Combinatorial hardness principle:** market-hard cases are stacked features (basket coupling + event memory + clocks + optionality), not isolated exotica.
9. **Observation-lens caution:** Markov/non-Markov appearance can be representation-dependent; structural non-commutation constraints are not.
10. **Evidence-first governance:** each promoted result requires reproducible artifacts (config, manifest, raw outputs, summary, command path).
11. **Scope/non-claim symmetry:** each strong claim must be paired with explicit failure boundaries and non-claims.
12. **Layer split discipline:** fin_spectral_time carries dynamics/computability; meta_theory_spectral_informa carries inferential state-object semantics.

These guidance points are the durable filter for deciding what counts as real progress versus narrative-only progress in this program.

6.4.5 Lane Closure Note: Structural Classes and Portability

To close the current lane explicitly, we record the core conceptual outcome:

1. Path dependence is treated as a structural class problem (direct / augmentable / impractical / hybrid), not as a product-name list.
2. The classing logic is partially method-independent: it is fundamentally about finite-state memory compressibility.
3. The class-to-solver mapping in this paper is method-specific: it determines when pure spectral, spectral-hybrid, or fallback lanes are appropriate.

What is unique in the present program is not each ingredient in isolation, but the integrated package:

- structural classing,
- operational decision protocol,
- claim-governance via explicit evidence gates,
- and reproducible artifact requirements.

This makes the framework portable beyond the spectral stack itself. Candidate reuse domains include:

- PDE/FEM solver triage for high-memory dynamics,
- Monte Carlo state-design and variance-control pipelines,
- control/RL state-construction for long-memory environments,
- model-risk governance with explicit non-claim boundaries.

Hence the lasting contribution is a computability-governance layer that can sit above multiple numerical engines, with spectral methods as one strongly compatible engine rather than the only possible backend.

6.5 General Path-Memory Framework (Beyond Finance)

The path-dependence classification above is not finance-specific. It is a general memory-computability framework for stochastic dynamical systems:

1. **Compressible memory**: history can be summarized by a finite-dimensional state augmentation
$$Z_t = (X_t, Y_t), \qquad Y_t = \mathcal{F}(X_{[0,t]}),$$
such that Z_t is Markov (or well-approximated as Markov).
2. **Non-compressible memory**: no finite-dimensional closure captures the required path information with acceptable error.

In the compressible class, spectral methods apply to the augmented generator $\mathcal{L}_Z$: memory is represented by additional state coordinates and propagated through the same semigroup/eigenmode machinery. In the non-compressible class, pure spectral closure is typically not sufficient and hybrid methods are required.

This distinction appears in multiple domains:

Domain	Typical memory mechanism	Spectral status
Exotic option pricing	running average, running extrema, occupation time, event history	often compressible (augmentable)
Hysteretic materials	loading-cycle history dependence	mixed: often low-dimensional internal-variable closure exists locally
Viscoelastic media	stress depends on deformation history (memory kernel)	often non-compressible unless kernel admits finite-state approximation
Anomalous diffusion/fractional dynamics	long-tail temporal memory	typically non-compressible in pure finite-state form
Open quantum systems	environment-induced non-Markov memory	mixed: Markovian embedding possible in restricted regimes

Hence, the core contribution is not only a pricing method but a computability lens: determine whether path memory is finite-state compressible, then choose pure spectral or hybrid architecture accordingly.

6.6 Living Systems as Path-Memory Dynamics

The same formal distinction applies to living systems. An organism is not determined by instantaneous physical state alone; it carries layered memory across multiple timescales:

- evolutionary memory (population-level selection history),
- developmental and physiological memory (organ-level internal-state history),
- cognitive and behavioral memory (beliefs, learned policies, context traces).

A minimal abstract model is again:
$$Z_t = (X_t, Y_t),$$

where X_t is observable external state and Y_t is internal memory state. The key modeling question is whether Y_t admits a finite-dimensional closure for the task under study.

If yes, one can use Markov embedding plus spectral propagation on the augmented generator. If no, one should treat the system as genuinely long-memory and use nonlocal or hybrid architectures.

This perspective does not reduce biology or cognition to finance; it identifies a shared mathematical structure: path dependence with varying degrees of memory compressibility.

6.7 Observation-Lens Principle

Markovianity versus path dependence is generally relative to the chosen state and observation map. If S_t is the full latent state and $O_t = h(S_t)$ is the observed state, then:

- S_t may be Markov while O_t is non-Markov (memory appears after projection),
- memory can disappear at a coarse level if the unresolved components are absorbed into effective noise,
- and memory can reappear when the state representation is refined.

So "non-path-dependent" is often an effective statement at a specific observational resolution, not an absolute ontological statement.

In quantum systems this distinction must be stated carefully: the Heisenberg uncertainty relation

$$\Delta A \, \Delta B \geq \frac{1}{2} \, |\langle [A, B] \rangle|$$

is a structural theorem for non-commuting observables within quantum mechanics, not a removable instrumentation artifact. What is lens-dependent is which observables, coarse-grainings, and effective state variables are chosen for modeling.

6.8 Bridge to the Spectral Information State and nested depth collapse

This paper's generator framework is a concrete instantiation of the broader state-object program in topics/meta_theory_spectral_information_state.

In the present setting, a natural time-state object is:

$$\mathcal{K}_t^{\mathrm{ST}} = (M_t^{\mathrm{aug}}, A_t, \mathcal{R}_t, \mathcal{E}_t),$$

where:

- M_t^{aug} is the base or augmented generator,
- A_t are current spectral coefficients,
- $\mathcal{R}_t$ is the readout family,
- $\mathcal{E}_t$ is the residual ledger.

From this single object, the main manifestations in this paper are readouts:

$$\text{American value}, \ \mathbb{E}[\tau_b], \ f_{\tau_b}, \ \psi(M), \ \text{and derivative prices} = \Phi_j(\mathcal{K}_t^{\mathrm{ST}}).$$

So the PDF/CDF layer is one manifestation, not the object itself.

This bridge also clarifies nested-expectation depth in XVA-style or multi-layer continuation settings. Write nested structure abstractly as:

$$V = b + T(V).$$

If T is a contraction ($\|T\| < 1$), then

$$V = \sum_{n=0}^{\infty} T^n b, \qquad \|V - V^{(L)}\| \le \frac{\|T\|^{L+1}}{1 - \|T\|} \|b\|.$$

Hence effective nesting depth is finite at fixed error tolerance.

Mode-wise, if $T\phi_k = \lambda_k \phi_k$, then:

$$|R_{k,L}| \le \frac{|b_k| \, |\lambda_k|^{L+1}}{1 - |\lambda_k|},$$

which makes depth planning a spectrum-controlled quantity. This is the same computational logic as the spectral-time thesis: temporal complexity is governed by eigen-structure, and residual accounting is part of the primary state object.

6.9 Notation alignment (shared with meta_theory_spectral_information_state)

To prevent symbol drift across manuscripts, we use:

- $\mathcal{K}_t$: primary stored object,
- $\mathcal{R}_t$: readout family,
- $\mathcal{E}_t$: residual ledger.

Within this paper:

- $\mathcal{K}_t^{\mathrm{ST}} = (M_t^{\mathrm{aug}}, A_t, \mathcal{R}_t, \mathcal{E}_t)$,
- M_t^{aug} carries generator structure,
- A_t carries coefficient-state information.

Eigenvalue symbols are disambiguated by superscripts:

- λ_k^M: generator-spectrum eigenvalues (intrinsic time scales),
- λ_k^T: nesting-operator eigenvalues in depth-collapse formulas.

This convention keeps the fixed-point nesting analysis and the generator-time analysis consistent without overloading one λ_k symbol for two distinct operators.

6.10 Durable Insights from the Path-Dependency Session

Source consolidation note: distilled from /.cursor/state/session−exports/cursor_path_dependency.md and promoted only where the idea has stable theorem/program value.

The following points are promoted from session discussion into stable research/program guidance:

1. **Computability-first rule:** first classify path dependence by memory compressibility, then choose the numerical solver.
2. **Finite-state closure criterion:** pure spectral deployment is strongest when path memory admits practical finite-dimensional Markov augmentation.
3. **Boundary honesty:** non-compressible memory is an explicit unsolved region for pure spectral closure; route these cases to spectral-hybrid pipelines.
4. **Jump unification:** Lévy/subordination effects remain in the same operator language through $M \mapsto \psi(M)$, with non-Markov clocks outside simple eigenvalue remapping.

5. **Combinatorial hardness:** hardest market products are stacked-feature structures (basket coupling + event memory + occupancy clocks + optionality), not isolated single-feature exotics.
6. **Evidence discipline:** claims must be supported by artifacts (config, manifest, raw outputs, summary, reproducible command surface), not narrative alone.
7. **Derivation-first theorem policy:** closed-form statements are promoted only after PO-chain discharge (PO-1..PO-7), not from symbolic pattern matching.
8. **Two-layer architecture:** fin_spectral_time is the dynamics/computability layer; meta_theory_spectral_information_state is the inferential/state-object layer.
9. **Observation-lens clarification:** Markov/non-Markov appearance can be representation-dependent, while structural non-commutation facts remain theorem-level constraints.
10. **Gate-based program control:** theorem substance, benchmark reality, boundary honesty, and formalization integrity gates are mandatory for promotion.

6.11 Session-Derived Open Research Questions

The following open questions are promoted from the session and kept as explicit research targets:

1. **Closure theorem sharpness:** give necessary-and-sufficient conditions for when a path functional admits practical finite-state Markov closure at a prescribed error budget.
2. **Hybrid optimality boundary:** characterize when spectral-hybrid strictly dominates pure spectral (runtime/error) for high-memory path products.
3. **Jump-plus-memory interaction:** formalize when subordinated-generator remapping $M \mapsto \psi(M)$ remains computationally stable under augmented path states.
4. **Second-regime CEV theorem:** complete the $\mu > 0,\ 0 < \beta < 1$ branch admissibility/uniqueness story or prove an impossibility boundary.
5. $\beta \to 1$ **asymptotic bridge:** derive and verify a dedicated rescaling theorem that connects the current CEV transform to the near-lognormal edge.
6. **Cross-domain transfer test:** quantify how the compressible/non-compressible memory lens transfers from exotic pricing to one non-finance domain without losing falsifiability.

Promotion rule for this question set:

- an item moves from “open question” to “resolved contribution” only with a theorem statement (or impossibility statement), reproducible evidence artifacts, and explicit boundary/non-claim language.

---

7. Limitations and Future Work

1. **Forward–backward distinction.** The Fokker–Planck matrix M evolves density forward; the backward (Kolmogorov) evolution for value functions requires $M^\top$. For processes with drift, $M \neq M^\top$ and using the wrong operator gives $O(1)$ errors. This is now correctly implemented; the American put matches MC within 0.8%.

2. **Multi-dimensional first passage.** The killed generator extends to d dimensions via tensor product of 1D killed and reflecting bases. The URRT controls the total spectral size. Numerical validation is needed.

3. **Non-Markovian subordinators.** The Bochner–Phillips theorem (Theorem 2) requires the subordinator to be Lévy. For non-Markovian time changes (e.g., fractional Brownian motion as clock), the simple eigenvalue remap $\psi(\lambda_k)$ does not apply; a more general Volterra-type operator is needed.
4. **Lean formalization.** The first passage formula (10) is straightforward to formalize in Lean 4 given the existing semigroup infrastructure in LeanProofs/SpectralFenton/. The subordination theorem (14) requires formalizing matrix functions and the Bochner–Phillips theorem.

---

8. Conclusion

Time in stochastic processes is an eigenvalue spectrum. This paper shows that three fundamental time-related problems — when to stop, when the process arrives, and what happens under random time change — all reduce to algebraic operations on the Fokker–Planck generator matrix:

$$\boxed{\begin{aligned} \text{Optimal stopping:}\quad & V(t) = \max\left(g,\ e^{M\Delta t}\, V(t+\Delta t)\right) \\ \text{First passage:}\quad & \mathbb{E}[\tau] = -\mathbf{1}^\top M_{\text{killed}}^{-1}\, A(0) \\ \text{Subordination:}\quad & \mathcal{L}_{\text{VG}} = \psi(M) = V\,\text{diag}(\psi(\lambda_k))\, V^{-1} \end{aligned}}$$

The generator M encodes all temporal structure. Clock time enters only through $e^{\lambda_k t}$. Subordination remaps λ_k. Barriers make all λ_k negative. Optimal stopping selects the time at which the mode sum exceeds a threshold. One matrix, every temporal phenomenon.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Bochner, S (1955). Harmonic Analysis and the Theory of Probability. *Harmonic Analysis and the Theory of Probability.*
- Carr, P., H. Geman, D. B. Madan, and M. Yor (2002). The fine structure of asset returns: An empirical investigation. *Journal of Business*, 75(2).
- Fang, Fang and Oosterlee, Cornelis W. (2008). A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions. *SIAM Journal on Scientific Computing*, 31(2), 826-848. DOI: 10.1137/080718061
- Longstaff, F.A., & Schwartz, E.S (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14(1), 113-147. DOI: 10.1093/rfs/14.1.113

- Madan, D. B., P. P. Carr, and E. C. Chang (1998). The Variance Gamma process and option pricing. *Review of Finance*, 2(1).
- Nagy, T. (2026). Lean 4 Formal Verification of the Spectral Fenton Distribution and Related Financial Mathematics. *Working paper.*
- Nagy, T. (2026). The Spectral Tensor Representation of Stochastic Processes. *Working paper.*
- Nagy, T. (2026). American Basket Option Pricing via Eigenvalue-Conditional COS Backward Induction: Convergence Rates and Complexity Analysis. *Working paper.*
- Nagy, T. (2026). Frequency-Domain Theory of Financial Economics: Thirteen Fundamental Results from One Decomposition. *Working paper.*
- Nagy, T. (2026). The Spectral Tensor Representation of Stochastic Processes. *Working paper.*
- Phillips, R. S (1952). On the generation of semigroups of linear operators. *Pacific Journal of Mathematics*, 2(3).

Appendix: Reproducibility

```
python3 examples/spectral_time.py
```

Self-contained, requires only NumPy and SciPy. Runtime: $ $80 seconds.