Spectral Interest Rate Pricing: COS-Based Derivatives from Yield Curve Coefficients

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Working Paper

Abstract

We present the first formally verified interest rate derivatives pricing engine. The Nagy spectral yield curve represents yields as a finite cosine series $y(\tau) = A_0 + \sum_{k=1}^{K} A_k \cos(k\pi\tau/\tau_{\max})$ with each mode following independent Ornstein-Uhlenbeck dynamics. When the COS method (Fang-Oosterlee, 2008) is applied to this spectral model, the density coefficients F_n are determined algebraically by the yield curve coefficients $\{A_k\}$ and their dynamics $\{\kappa_k, \sigma_k\}$, giving caplet prices as finite cosine sums with exponential convergence $O(\rho^{-N})$. We prove in Lean 4: (a) COS caplet pricing converges exponentially for analytic densities; (b) cap-floor parity holds exactly in the COS framework; (c) spectral Greeks $\partial\text{Caplet}/\partial A_k$ have the same COS structure as pricing, enabling analytical computation via the chain rule through the OU characteristic function; (d) pricing is deterministic by construction. The complete framework — discount factors, forward LIBOR, COS caplet/cap pricing, Jamshidian swaption decomposition, and spectral Greeks — comprises 10 Lean files with 66 statements (zero sorry), importing 4 cross-domain theorems from 3 previously verified libraries. A companion Rust implementation validates the framework in a lognormal special case (see Section 4.4 for scope of verification).

---

1. Introduction

1.1 The Problem

Interest rate derivatives — caplets, caps, floors, swaptions — constitute one of the largest segments of global derivatives markets. Pricing these instruments requires: (i) a term structure model specifying yield curve dynamics; (ii) a numerical method for computing expectations under the risk-neutral measure; (iii) sensitivity computations (Greeks) for hedging. The standard approaches — tree-based methods, Monte Carlo simulation, PDE solvers — are well-understood but share a common deficiency: none provides formal guarantees of correctness. Pricing code is tested, not proved.

This matters because interest rate models are compositional. A cap is a portfolio of caplets; a swaption decomposes (under Jamshidian's trick) into a portfolio of bond options; Greeks are derivatives of prices with respect to model parameters. Each composition step introduces potential errors: truncation, discretization, numerical differentiation. The interactions between these approximations are tested empirically but not verified formally.

1.2 The Spectral Yield Curve

The Nagy spectral yield curve (Nagy, 2026; building on Nelson-Siegel, 1987, and Diebold-Li, 2006) represents yields as a cosine series:

$$y(\tau) = A_0 + \sum_{k=1}^{K} A_k \cos\left(\frac{k\pi\tau}{\tau_{\max}}\right)$$

Each mode follows independent Ornstein-Uhlenbeck dynamics:

$$dA_k(t) = -\kappa_k(A_k(t) - \bar{A}_k)\,dt + \sigma_k\,dW_k(t)$$

This model has been previously verified in Lean 4: the NagyYieldCurve structure in SpectralTrading/NagyYieldCurve.lean proves Fejér-based no-arbitrage, $O(K)$ verification of coefficient bounds, and that the spectral model subsumes Nelson-Siegel as a 3-mode special case.

1.3 The COS Method Meets Spectral Yield Curves

The COS method (Fang and Oosterlee, 2008) prices European-style derivatives by expanding the risk-neutral density in a cosine series on $[a, b]$:

$$\mathbb{E}[\max(L - K, 0)] = \sum_{n=0}^{N-1}{}' F_n \cdot V_n$$

where $F_n = \frac{2}{b-a}\,\mathrm{Re}\big[\varphi\big(\frac{n\pi}{b-a}\big)\,e^{-in\pi a/(b-a)}\big]$ are density coefficients from the characteristic function φ, and V_n are payoff coefficients (closed-form for call payoffs).

The key insight: when the underlying model IS already a cosine series, the COS expansion becomes algebraic. The characteristic function φ of the forward LIBOR rate under the T_2-forward measure is determined directly by $\{A_k, \kappa_k, \sigma_k\}$. No simulation, no PDE, no trees.

1.4 Our Contribution

We formally verify, in Lean 4, that this spectral COS pricing framework is:

1. **Convergent** at rate $O(\rho^{-N})$ for analytic densities (exponential, not polynomial);
2. **Arbitrage-free**: cap-floor parity holds exactly within the COS framework;
3. **Analytically differentiable**: spectral Greeks $\partial\text{Caplet}/\partial A_k$ have the same COS structure as the price, so the derivative of the sum equals the sum of derivatives — given an analytical expression for $\partial F_n/\partial A_k$, no numerical differentiation is required;
4. **Deterministic**: prices are finite arithmetic expressions — no simulation variance.

The verification consists of 10 Lean 4 files organized in a 4-tier dependency DAG, with zero sorry (unproved assertions). Four cross-domain imports connect this work to previously verified results on COS truncation error, mixture collapse linearity, and dimension-free convergence.

1.5 Prior Art

No formally verified interest rate derivatives pricing engine exists in the literature. The closest related work includes:

- **Lean-verified financial mathematics**: Harrison-Pliska fundamental theorems in Lean [TODO:cite Avigad et al.], Black-Scholes in Isabelle/HOL [TODO:cite Hölzl et al.]. These verify individual theorems, not pricing engines.
- **COS method theory**: Fang and Oosterlee (2008, 2009) establish convergence for European and Bermudan options. Our contribution is the formal verification of convergence in the spectral setting, with explicit cross-domain imports.
- **Spectral yield curve models**: Nelson-Siegel (1987), Diebold-Li (2006), Svensson (1994). Our model extends these to a general cosine basis with verified no-arbitrage.
- **COS method extensions**: Ruijter and Oosterlee (2012) extend COS to two dimensions [TODO:cite Ruijter & Oosterlee 2012, SIAM SISC]; Junike and Pankrashkin (2022) sharpen COS convergence bounds [TODO:cite Junike & Pankrashkin 2022]. Our work applies COS in the specific setting of spectral yield curves and adds formal verification.
- **Affine term structure models**: Duffie and Kan (1996) and Dai and Singleton (2000) provide analytically tractable term structure models with known characteristic functions [TODO:cite Duffie & Kan 1996; Dai & Singleton 2000]. Our spectral OU model is a special case of the affine class; the contribution is the cosine-basis structure enabling COS algebraicity and the formal verification layer.

---

2. Preliminaries

2.1 Notation

Symbol	Meaning
$y(\tau)$	Yield at maturity τ
A_k	k-th spectral coefficient of the yield curve
κ_k, σ_k	Mean-reversion speed and volatility of mode k
$P(0,\tau)$	Discount factor: $\exp(-y(\tau)\cdot\tau)$
$L(T_1,T_2)$	Forward LIBOR rate over $[T_1,T_2]$
δ	Accrual period: $T_2 - T_1$
F_n	COS density coefficient (from characteristic function)
V_n	COS payoff coefficient (from call payoff integral)
ρ	Bernstein ellipse parameter: $\rho > 1$ for analytic densities
N	Number of COS expansion terms
S	Par swap rate

2.2 Discount Factor

Definition. For yield y at maturity τ:

$$P(0, \tau) = \exp(-y(\tau) \cdot \tau)$$

This maps any real yield to a strictly positive discount factor.

2.3 Forward LIBOR

Definition. The simple forward rate over $[T_1, T_2]$:

$$L(T_1, T_2) = \frac{1}{\delta} \left(\frac{P(0, T_1)}{P(0, T_2)} - 1 \right)$$

where $\delta = T_2 - T_1$. Well-defined since $P(0, T_2) > 0$.

2.4 COS Density Coefficients

Definition. The n-th density coefficient:

$$F_n = \frac{2}{b-a} \operatorname{Re}\left[\varphi\left(\frac{n\pi}{b-a} \right) \exp\left(-\frac{in\pi a}{b-a} \right) \right]$$

For distributions with characteristic functions analytic in a strip of width $\log \rho$, $|F_n| = O(\rho^{-n})$ with $\rho > 1$.

2.5 COS Payoff Coefficients

Definition. For the call payoff $\max(x - K, 0)$ on $[a, b]$:

$$V_n = \frac{1}{b-a} \int_K^b (x - K) \cos\left(\frac{n\pi(x-a)}{b-a} \right) dx$$

The zeroth coefficient $V_0 = (b-K)^2/(2(b-a)) \geq 0$. Higher-order coefficients decay as $O(1/n^2)$.

2.6 Characteristic Function and Domain Selection

The characteristic function. This section fills the gap between the spectral yield curve model (Section 1.2) and the COS pricing engine (Section 1.3).

Each spectral mode A_k follows OU dynamics (Section 1.2). Conditional on time-0 values, each $A_k(T)$ is Gaussian:

$$A_k(T) \sim \mathcal{N}\left(\bar{A}_k + (A_k(0) - \bar{A}_k) e^{-\kappa_k T},\ \frac{\sigma_k^2}{2\kappa_k}(1 - e^{-2\kappa_k T}) \right)$$

The zero-coupon bond price at maturity T is $P(0, T) = \exp(-y(T) \cdot T)$ where $y(T) = A_0 + \sum_{k=1}^{K} A_k \cos(k\pi T/\tau_{\max})$. The forward LIBOR rate $L(T_1, T_2)$ is a ratio of bond prices (Section 2.3). The COS method is applied to $x = \log L(T_1, T_2)$, the log-forward rate, because the logarithm of a ratio of log-normal quantities has a tractable characteristic function.

For the single-mode ($K = 0$) lognormal special case — which the Rust implementation uses — the forward rate L is lognormal and the characteristic function of $x = \log L$ is:

$$\varphi(u) = \exp\left(iu\mu_x - \frac{u^2\sigma_x^2}{2}\right)$$

where $\mu_x = \log L(T_1, T_2) - \sigma_x^2/2$ (drift-adjusted log mean) and σ_x^2 is the variance of the log-forward rate over $[0, T_1]$. For the general multi-mode case ($K \geq 1$), the log-forward rate is a nonlinear function of $K + 1$ independent Gaussians $\{A_k(T)\}$, and the characteristic function does not have a closed form. In this case, one may use the Fenton-Wilkinson lognormal approximation (Nagy, 2026) or moment-matching to obtain an approximate $\varphi(u)$; the COS convergence guarantee holds provided the approximated density is analytic.

Domain selection. The COS expansion is computed on a finite domain $[a, b]$ for $x = \log L$. The domain must be wide enough to capture the tails of the density but not so wide that the oscillatory cosine basis wastes resolution. The Rust implementation uses the cumulant-based rule:

$$a = \mu_x - 10\sigma_x, \qquad b = \mu_x + 10\sigma_x$$

where μ_x and σ_x are the mean and standard deviation of $x = \log L$. The factor of 10 ensures that the domain truncation error $\varepsilon_{\text{domain}} = O(e^{-c(b-a)^2})$ is negligible relative to the COS truncation error for any practical N. For heavier-tailed models, the factor may need to increase; Fang and Oosterlee (2008) recommend a similar cumulant-based rule with a factor of $L = 10$–12.

---

3. Main Results

3.1 Theorem 1 — Discount Factor Positivity

[Lean-verified, DiscountFactor.lean, 0 sorry]

Theorem. *For any yield $y \in \mathbb{R}$ and maturity $\tau \geq 0$:*

(a) $P(0, \tau) > 0$ (strict positivity);

(b) $P(0, 0) = 1$ (unit at zero maturity);

(c) $P(0, \tau) \leq 1$ when $y \geq 0$ and $\tau \geq 0$;

(d) P is decreasing in y for $\tau \geq 0$ (higher yield → lower price);

(e) P is decreasing in τ for $y \geq 0$ (longer maturity → lower price).

Proof architecture. All properties follow from monotonicity of exp: Real.exp_pos for (a), Real.exp_zero for (b), Real.exp_le_exp.mpr with nlinarith for (c)-(e). The proof imports the NagyYieldCurve structure from SpectralTrading/ for the yield curve model definition.

Cross-import: LeanProofs.SpectralTrading.NagyYieldCurve — yield curve structure, Fejér no-arb bound.

3.2 Theorem 2 — Forward LIBOR Positivity

[Lean-verified, ForwardLIBOR.lean, 0 sorry]

Theorem. *If $y_2{\cdot}T_2 > y_1{\cdot}T_1$ (the standard upward-sloping yield curve condition) and $\delta = T_2 - T_1 > 0$, then $L(T_1, T_2) > 0$.*

Proof architecture. The condition $y_2 T_2 > y_1 T_1$ implies $P(0, T_2) < P(0, T_1)$ via Real.exp_lt_exp.mpr, which gives $P(0, T_1)/P(0, T_2) > 1$, whence $(P_1/P_2 - 1)/\delta > 0$ by div_pos. The well-definedness of division is guaranteed by Theorem 1(a): $P(0, T_2) > 0$.

Remark (inverted yield curves). The hypothesis $y_2 T_2 > y_1 T_1$ is equivalent to requiring that the zero-coupon curve is upward-sloping in the "dollar-duration" sense over $[T_1, T_2]$. This fails for inverted yield curves, which do occur in practice (e.g., pre-recession environments). When the curve is inverted over the caplet's accrual period, the forward rate $L(T_1, T_2) < 0$ in the model, and caplet pricing requires careful treatment of the negative-rate regime. The current framework does not address negative rates; extending it (e.g., via shifted lognormal or Bachelier dynamics) is left as future work.

3.3 Theorem 3 — COS Density Coefficient Decay

[Lean-verified, COSDensityCoefficients.lean, 0 sorry]

Theorem. *For distributions with characteristic functions analytic in a strip of width $\log \rho$ ($\rho > 1$):*

(a) $|F_n| \le C_\varphi \cdot |\varphi(n\pi/(b - a))|$ where $C_\varphi = 2/(b - a)$;

(b) $|F_n| \le C \cdot \rho^{-n}$ (exponential decay);

(c) The COS truncation error satisfies $\varepsilon_N \le C\rho^{-N}/(1 - \rho^{-1})$ (geometric tail sum);

(d) The COS truncation is the first component of the 6-part error decomposition:

$$|\varepsilon_{\text{total}}| \le \varepsilon_N + \varepsilon_{\text{GH}} + \varepsilon_{\text{outer}} + \varepsilon_{\text{domain}} + \varepsilon_{\text{res}} + \varepsilon_{\text{fp}}$$

Proof architecture. Part (c) is the genuine content: the truncation error bound is derived from the geometric tail sum of exponentially decaying coefficients, using cos_truncation_error_bound with positivity of the geometric series. Parts (a) and (b) are structural: the Lean statements assume the density coefficient bound and exponential decay as hypotheses and propagate them through the framework. Deriving $|F_n| \le (2/(b-a)) \cdot |\varphi(n\pi/(b-a))|$ from the definition of F_n (using $\text{Re}(z) \le |z|$, which is Complex.re_le_abs in Mathlib) is straightforward but is left to the hypothesis level in the current formalization. Part (d) is imported from the Spectral Fenton framework.

Cross-import: LeanProofs.SpectralFenton.ErrorDecomposition — the 6-component error decomposition theorem error_decomposition_six, establishing the triangle inequality structure.

3.4 Theorem 4 — Payoff Coefficient Properties

[Lean-verified, PayoffCoefficients .lean, 0 sorry]

Theorem. *For the call payoff $\max(x - K, 0)$ on $[a, b]$ with $a < b$ and $K \le b$:*

(a) $V_0 \ge 0$ (non-negative expected payoff);