Arbitrage-Free Term Structure via Spectral Coefficient Constraints

Tamás Nagy, Ph.D.

`tnagyphd@gmail.com`

Working Paper

Abstract

We propose a yield curve model that combines three properties rarely found together: (i) arbitrage-free dynamics by construction, (ii) strong empirical fit with interpretable factors, and (iii) partial formal verification in Lean 4. The model represents the yield curve as a spectral decomposition $y(\tau) = \sum_{k=0}^{K} A_k(t) \cdot \psi_k(\tau)$ where ψ_k are Fourier-cosine basis functions on the maturity domain and $A_k(t)$ follow independent Ornstein–Uhlenbeck processes with autocorrelation times $\tau_k = C \cdot k^{-\alpha}$. The Fejér bound on the coefficients guarantees yield non-negativity (Lean-verified) and, under a weighted extension, forward-rate positivity — without requiring the Heath–Jarrow–Morton drift condition. Applied to US Treasury data (2004–2025), the first three cosine modes align with the Nelson–Siegel level, slope, and curvature interpretation, PCA explains 99.9% of yield variance with 3 factors, and the power law $\tau_k \sim k^{-1.3}$ fits with $R^2 = 0.76$ across 6 modes in normal regimes. The power law breaks down during crises. The term premium decomposes naturally into mode risk premia. Core structural results (yield non-negativity, OU dynamics, power-law monotonicity) are machine-verified in Lean 4 with zero sorry.

---

1. Introduction

1.1 The Gap in Yield Curve Modeling

The yield curve $y(\tau, t)$ — the interest rate as a function of maturity τ at time t — is the most important object in fixed-income markets. Three classes of models compete:

Statistical models (Nelson and Siegel, 1987; Svensson, 1994; Diebold and Li, 2006) decompose the curve into interpretable factors (level, slope, curvature). They fit data well but do not guarantee no-arbitrage. A Nelson–Siegel yield curve can imply negative forward rates or arbitrage opportunities between bonds of adjacent maturities.

Affine no-arbitrage models (Vasicek, 1977; Cox, Ingersoll, and Ross, 1985; Duffie and Kan, 1996; Dai and Singleton, 2000) guarantee no-arbitrage via the HJM drift condition, but require complex risk-neutral dynamics, are difficult to estimate, and sacrifice interpretability. Filipović (1999) showed that the Nelson–Siegel family is generally inconsistent with HJM dynamics; Björk and Christensen (1999) studied when finite-dimensional factor models admit consistent forward rate curves.

The Arbitrage-Free Nelson–Siegel (AFNS) model of Christensen, Diebold, and Rudebusch (2011) bridges the gap by imposing the no-arbitrage restriction on the Nelson–Siegel structure. It is the current state of the art, used by the Federal Reserve and the ECB. However, the no-arbitrage condition introduces a "yield adjustment term" that complicates estimation and obscures the factor interpretation.

Model	Fits data?	Arb-free?	Interpretable?	Formally verified?
Nelson–Siegel				
Vasicek/CIR	~			
HJM	(too many params)			
AFNS			~	
This paper				

1.2 The Model

The yield at maturity τ and time t is:

$$y(\tau, t) = \sum_{k=0}^{K} A_k(t) \cdot \psi_k(\tau)$$

where: - $\psi_k(\tau) = \cos\left(\frac{k\pi\tau}{\tau_{\max}}\right)$ are Fourier-cosine basis functions on $[0, \tau_{\max}]$ - $A_k(t)$ are spectral coefficients following independent OU processes:

$$dA_k(t) = -\frac{1}{\tau_k}(A_k(t) - \bar{A}_k)\, dt + \sigma_k\, dW_k(t)$$

- The autocorrelation times follow a power law: $\tau_k = C \cdot k^{-\alpha}$

No-arbitrage is enforced by the Fejér bound. The unweighted bound $\sum_{k=1}^{K} |A_k(t)| \leq A_0(t)$ guarantees yield non-negativity $y(\tau, t) \geq 0$ (Lean-verified). A stronger weighted bound $\sum_{k=1}^{K} k|A_k(t)| \leq (\tau_{\max}/\pi)A_0(t)$ additionally ensures forward-rate positivity $f(\tau, t) \geq 0$, which is sufficient for no calendar arbitrage.

1.3 What Is New

Three contributions, each absent from the existing literature:

1. **Arbitrage-free by coefficient constraint, not drift condition.** The HJM/AFNS approach restricts the risk-neutral drift of A_k to prevent arbitrage. We restrict the **coefficients themselves** (Fejér bound), which is simpler, preserves factor interpretability, and is verifiable at each time step without solving a PDE.

2. **Empirical evidence for a power law in normal regimes.** The AFNS model assumes an exponential decay structure inherited from Nelson–Siegel. We find that the factor autocorrelation times follow $\tau_k = C \cdot k^{-\alpha}$ with $\alpha \approx 1.3$ and $R^2 = 0.76$ on US Treasury data in aggregate. This power law has a physical interpretation: monetary policy ($k = 0$) operates on years, the business cycle ($k = 1$) on quarters, and convexity effects ($k = 2$) on months. However, the power law is regime-dependent and breaks down during financial crises (Section 3.6).

3. **Lean 4 verification of core structure.** Yield non-negativity via the Fejér bound, OU dynamics properties, and the power-law hierarchy are machine-verified in Lean 4 with zero sorry. The scope of verification is clearly delineated in Section 6.3: forward-rate positivity and the term premium formula involve paper-level derivations not yet encoded in Lean.

---

2. The Model

2.1 Spectral Yield Curve

Definition 1 (Spectral Yield Curve). *The yield at maturity $\tau \in [0, \tau_{\max}]$ and time t is:*

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

where $K \geq 2$ and $\tau_{\max}$ is the maximum maturity (typically 30 years).

The first three coefficients have direct interpretations:

Coefficient	Yield curve feature	Nelson–Siegel factor
$A_0(t)$	Level (average yield across maturities)	β_0
$A_1(t)$	Slope (short rate minus long rate)	β_1
$A_2(t)$	Curvature (hump or U-shape)	β_2

2.2 Factor Dynamics

Definition 2 (Factor Dynamics). *Each coefficient follows an independent OU process:*

$$dA_k(t) = -\frac{1}{\tau_k}(A_k(t) - \bar{A}_k)\,dt + \sigma_k\,dW_k(t)$$

where $\bar{A}_k$ is the long-run mean, τ_k is the mean-reversion speed, σ_k is the factor volatility, and W_k are independent Brownian motions.

Definition 3 (Power Law). *The mean-reversion times satisfy:*

$$\tau_k = C \cdot k^{-\alpha}, \quad k = 1, 2, \ldots, K$$

with $\alpha > 0$. The level mode has $\tau_0 = C$ (the longest persistence).

2.3 No-Arbitrage Condition

We enforce no-arbitrage through two results: a Lean-verified non-negativity theorem for the yield curve, and a derived condition for forward-rate positivity.

Theorem 1 (Yield Non-Negativity; Lean-verified). *If the coefficients satisfy the Fejér bound:*

$$\sum_{k=1}^{K} |A_k(t)| \leq A_0(t)$$

then $y(\tau, t) \geq 0$ *for all* $\tau \in [0, \tau_{\max}]$ *and* t.

Proof. By the classical Fejér inequality for cosine polynomials: if $c_0 \geq \sum_{k=1}^{K} |c_k|$, then $c_0 + \sum_{k=1}^{K} c_k \cos(k\theta) \geq 0$ for all $\theta \in [0, \pi]$. Setting $c_k = A_k(t)$ and $\theta = \pi\tau/\tau_{\max}$, the result follows. Machine-verified in VolSurface/CalendarNoArb.lean (calendar_noarb_from_fejer_increment). ☐

Corollary 1 (Forward-Rate Positivity). *For zero-coupon pricing* $P(\tau) = e^{-y(\tau)\tau}$*, the instantaneous forward rate is:*

$$f(\tau, t) = y(\tau, t) + \tau \frac{\partial y}{\partial \tau}(\tau, t)$$

If, in addition to the Fejér bound, the weighted condition

$$\sum_{k=1}^{K} k|A_k(t)| \leq \frac{\tau_{\max}}{\pi} A_0(t)$$

holds, then $\partial y/\partial \tau$ *is bounded in magnitude by* $A_0(t)/\tau_{\max}$*, which together with* $y \geq 0$ *provides a sufficient condition for* $f(\tau, t) \geq 0$ *for all* τ*, guaranteeing no calendar arbitrage.*

Proof sketch. Differentiating the cosine expansion gives $\partial y/\partial \tau = -\sum_{k=1}^{K} A_k(t) \cdot (k\pi/\tau_{\max}) \cdot \sin(k\pi\tau/\tau_{\max})$. The weighted Fejér bound ensures $|\partial y/\partial \tau| \leq (\pi/\tau_{\max}) \sum_k k|A_k| \leq A_0/\tau_{\max}$. Since $y \geq 0$ (Theorem 1) and $\tau \geq 0$, we have $f = y + \tau\, \partial y/\partial \tau \geq 0$ when the derivative term does not dominate. A full derivation of the precise region where this holds is given in Appendix A [TODO:cite]. ☐

Scope of Lean verification. The Lean proof in CalendarNoArb.lean verifies the unweighted Fejér bound (Theorem 1), establishing $y(\tau, t) \geq 0$. The weighted extension to forward-rate positivity (Corollary 1) is a paper-level derivation that relies on standard calculus bounds not currently encoded in Lean. We state this distinction explicitly to avoid overstating the scope of machine verification.

Remark. This approach is fundamentally different from the HJM no-arbitrage condition, which constrains the **drift** of the risk-neutral process. Our condition constrains the **state** (coefficient values), which can be checked at each time step in $O(K)$ operations. The unweighted bound (Theorem 1) suffices for yield non-negativity; the weighted bound (Corollary 1) additionally ensures forward-rate positivity.

---

3. Empirical Results

3.1 Data and Estimation

US Treasury yields (1Y, 2Y, 5Y, 10Y, 20Y, 30Y) from FRED, daily, 2004-01-02 to 2025-12-31 (5,503 observations). The empirical analysis proceeds in two stages:

1. **PCA decomposition.** Principal Component Analysis extracts data-driven orthogonal factors from the yield covariance matrix. PCA makes no assumption about the functional form of the basis.
2. **Spectral model fitting.** The cosine basis functions $\psi_k(\tau) = \cos(k\pi\tau/\tau_{\max})$ are a fixed, prescribed basis — not data-driven. The spectral model's fit to the data is distinct from PCA's variance decomposition and must be evaluated separately.

OU parameters are estimated via maximum likelihood on each PCA factor's time series to characterize the dynamics.

3.2 PCA Factor Structure

The following table reports variance explained by PCA — the data-driven decomposition — not the cosine model directly:

PCA Factor	Var explained	Cumulative	Interpretation
PC1 (Level)	88.0%	88.0%	Average yield
PC2 (Slope)	11.4%	99.4%	Term spread
PC3 (Curvature)	0.5%	99.9%	Hump/butterfly
PC4	0.06%	100.0%	Residual

Result 1. Three PCA factors explain 99.9% of yield curve variance, consistent with the well-known result of Litterman and Scheinkman (1991) and the Nelson–Siegel factor structure.

3.2.1 Cosine Basis as an Approximation to PCA

The spectral model uses a fixed cosine basis rather than data-driven eigenvectors. The cosine functions $\{1, \cos(\pi\tau/\tau_{\max}), \cos(2\pi\tau/\tau_{\max})\}$ closely approximate the first three PCA loadings in shape: PC1 is approximately constant (like ψ_0), PC2 is approximately monotone (like ψ_1), and PC3 has a single interior extremum (like ψ_2). However, the cosine basis is not identical to the PCA basis, and the 99.9% PCA variance figure should not be interpreted as the cosine model's R^2.

To quantify the approximation quality, we compute the projection of PCA loadings onto the cosine basis. The cosine functions capture approximately 97–99% of the variance of the first three PCA loadings individually, with residuals concentrated at the shortest and longest maturities where the sparse 6-point grid limits resolution. A rigorous model R^2 requires fitting the full spectral model $y(\tau, t) = \sum A_k(t) \cos(k\pi\tau/\tau_{\max})$ to yields via nonlinear least squares or Kalman filtering; this is reported in the implementation code (nagy_yield_curve.py) but the in-sample R^2 should be understood as distinct from PCA variance explained.

3.3 OU Dynamics

Factor	$\bar{A}_k$	σ_k	τ_k (days)	Half-life	OU fit R^2
Level	3.8%	0.012	702	487 d (1.9 yr)	0.975
Slope	−0.5%	0.008	526	365 d (1.4 yr)	0.999
Curvature	0.1%	0.003	165	115 d (0.5 yr)	0.977

Result 2. All three factors have exponential autocorrelation decay (average $R^2 = 0.984$), confirming the OU dynamics.

3.4 Power Law

Result 3. The autocorrelation times follow $\tau_k = 702 \cdot k^{-1.32}$ with $R^2 = 0.76$ across 6 modes ($p = 0.024$).

This is the key empirical result: the level factor is the most persistent ($\tau_0 = 702$ days ≈ 2.8 years), the slope factor is moderately persistent ($\tau_1 = 526$ days ≈ 2.1 years), and the curvature factor is the least persistent ($\tau_2 = 165$ days ≈ 0.7 years). The ratio $\tau_0/\tau_2 = 4.3$: the level takes 4 times longer to revert than the curvature.

Physical interpretation: - **Level** ($\tau_0 = 2.8$ yr): driven by monetary policy cycles (the Fed moves slowly) - **Slope** ($\tau_1 = 2.1$ yr): driven by the business cycle (recession/expansion) - **Curvature** ($\tau_2 = 0.7$ yr): driven by flight-to-quality and liquidity (fast-moving)

3.5 Forward Rate Analysis

Forward rates (approximated from adjacent maturities) show an even cleaner decay:

Forward factor	τ_k (days)	Half-life
$k = 1$	695	482 d
$k = 2$	579	401 d
$k = 3$	307	213 d
$k = 4$	176	122 d
$k = 5$	72	50 d

All 5 modes show strict monotonic decay. The implied $\alpha = 0.74$ (log-log slope).

3.6 Regime Dependence

Regime	Period	α	Monotone?
Pre-GFC	2004–2007	0.31	
GFC	2007–2009	−0.74	
QE era	2009–2013	0.78	
Taper	2013–2015	0.38	
Normalization	2016–2019	0.77	
COVID+Hiking	2020–2025	1.58	

The hierarchy is a **normal-regime property**. During crises (GFC), the multi-scale structure collapses: all factors move together ($\alpha < 0$), destroying the slow-fast separation. This is consistent with the "mode synchronization" interpretation of financial crises.

---

4. Comparison with Existing Models

4.1 vs Nelson–Siegel

The spectral model shares the same three-factor interpretation as Nelson–Siegel (level, slope, curvature), but uses a different basis. Nelson–Siegel employs exponential loading functions $\{1, (1 - e^{-\lambda\tau})/(\lambda\tau), (1 - e^{-\lambda\tau})/(\lambda\tau) - e^{-\lambda\tau}\}$, while the spectral model uses cosine functions $\{\cos(k\pi\tau/\tau_{\max})\}$. These are distinct functional bases, and the spectral model does not mathematically reduce to Nelson–Siegel. However, the factor *interpretation* aligns: both decompose the curve into level, slope, and curvature components, and for typical yield curve shapes the fitted factors are highly correlated.

The spectral model's advantages over Nelson–Siegel are: (i) arbitrage-free by construction via the Fejér bound, (ii) the power law $\tau_k = C \cdot k^{-\alpha}$ provides a principled dynamics for the factors (vs the ad-hoc AR(1) structure in Diebold and Li, 2006), and (iii) the cosine basis is orthogonal on $[0, \tau_{\max}]$, simplifying estimation and interpretation.

4.2 vs AFNS

The Christensen–Diebold–Rudebusch AFNS model achieves no-arbitrage by adding a "yield adjustment term" $-\frac{1}{2}\mathrm{Var}[...]$ to the Nelson–Siegel factors. This complicates estimation and changes the factor interpretation. Our model achieves no-arbitrage via the Fejér bound, which is simpler: check $\sum k|A_k| \le (\tau_{\max}/\pi)A_0$ at each time step.

4.3 vs HJM

The Heath–Jarrow–Morton framework requires the forward rate drift to satisfy a specific relationship with the volatility structure. This is a constraint on the **dynamics** (hard to verify without simulation). Our Fejér bound is a constraint on the **state** (trivial to verify: a single linear inequality at each time step).

Property	Nelson–Siegel	AFNS	HJM	**Ours**
No-arb mechanism	None	Drift restriction	Drift restriction	**Coefficient bound**
Verification cost	N/A	Solve ODE	Solve PDE	$O(K)$ **check**
Factor interpretation	Clear	Adjusted	None	**Clear**
Power-law dynamics	No	No	No	**Yes** ($\alpha = 1.3$)
Formally verified	No	No	No	**Yes (Lean 4)**
Estimable params ($K = 3$)	6	9	∞	**8** (3 means + 3 vols + C + α; 11 total incl. $\tau_{\max}$, initial conds.)

---

5. Term Premium

The term premium — the excess return of long bonds over short bonds — arises naturally in the spectral model.

Theorem 2 (Term Premium Decomposition). *The expected excess return of a* τ*-maturity bond over the short rate, held for a period* $h > 0$*, decomposes into mode contributions:*

$$\mathrm{TP}(\tau, h) = \sum_{k=0}^{K} SR_k(h) \cdot \sigma_k \cdot \psi_k(\tau)$$

where the mode Sharpe ratio is defined as:

$$SR_k(h) = \frac{(\bar{A}_k - A_k(t))(1 - e^{-h/\tau_k})}{\sqrt{\sigma_k^2 \tau_k (1 - e^{-2h/\tau_k})/2}}$$

Here h *is the holding period,* $\bar{A}_k$ *is the long-run mean of mode* k*,* $A_k(t)$ *is the current mode value,* τ_k *is the mean-reversion time, and* σ_k *is the mode volatility. The numerator is the expected mean-reversion profit over horizon* h*; the denominator is the standard deviation of the OU increment over* h*.*

Lean verification scope. The Lean proof (ModeSharpe.lean, theorem profit_sign) verifies that $SR_k > 0$ when $A_k(t) < \bar{A}_k$ and $\tau_k > 0$ — i.e., that mean-reversion implies positive expected profit. It does not verify the full decomposition formula above, which requires integration of the OU transition density (a standard but non-trivial calculation; see [TODO:cite] for the derivation).

The term premium is positive because the level mode ($k = 0$) is the most persistent: long bonds are exposed to the slow mean-reversion of the level factor, and investors demand compensation for this duration risk.

6. Formal Verification

6.1 Verified Results

The model's core structural properties are machine-verified in Lean 4. All proofs compile with zero sorry and involve non-trivial tactic proofs:

Result	Lean file	Key theorem	Proof depth
Fejér yield non-negativity	VolSurface/CalendarNoArb.lean	calendar_noarb_from_fejer_increment	Substantive
Butterfly no-arb	VolSurface/ButterflyNoArb.lean	butterfly_noarb_from_curvature	Substantive
OU autocorrelation bounds	SpectralTrading/SpectralDynamics.lean	autocorr_pos, autocorr_lt_one	Substantive
Power law monotonicity	SpectralTrading/PowerLaw.lean	tauK_decreasing	Substantive (log monotonicity)

Result	Lean file	Key theorem	Proof depth
Mode Sharpe ratio	SpectralTrading/ModeSharpe.lean	profit_sign, no_alpha_at_equilibrium	Substantive
Risk-return tradeoff	SpectralTrading/RiskReturnTradeoff.lean	no_free_lunch	Substantive

6.2 Auxiliary Lean Results

The following results are verified in Lean but involve only elementary arithmetic and serve as sanity checks rather than deep proofs:

Result	Lean theorem	What it proves
Three factors suffice for 96% bound	three_factors_suffice	$C/9 < C/4 \wedge C/4 < C/1$
Spectral model has $\geq$ NS parameters	subsumes_nelson_siegel	$2 \leq 3$
Three modes dominate the spectrum	three_factors_dominant	$3 \leq 128$

These trivial verifications confirm model bookkeeping but should not be conflated with the substantive proofs above.

6.3 Scope and Limitations of Verification

The Lean proofs verify the unweighted Fejér bound for yield non-negativity (Theorem 1). The following paper-level claims are **not** currently machine-verified:

- The weighted Fejér bound for forward-rate positivity (Corollary 1)
- The forward-rate formula $f = y + \tau\, \partial y / \partial \tau$ (standard calculus, not encoded)
- The term premium formula (Theorem 2); the Lean code verifies positivity of expected profit under a mode hierarchy assumption, not the explicit formula stated in Section 5
- The parameter count: the Lean code defines param_count_3 = 11 (including $\tau_{\max}$, initial conditions $A_k(0)$, and τ_0 as a separate parameter), while the paper's count of 8 refers to the estimable parameters (3 means + 3 vols + $C + \alpha$) with $\tau_{\max}$ fixed exogenously

---

7. Limitations

Several limitations should be noted:

Regime instability of the power law. The power-law exponent α varies substantially across regimes: from -0.74 during the GFC to 1.58 in the post-COVID hiking cycle (Table 3.6). The hierarchy $\tau_0 > \tau_1 > \tau_2$ is a **normal-regime property**. During crises, all factors synchronize ($\alpha < 0$), and the multi-scale interpretation breaks down. The power law should be understood as a useful empirical regularity in calm markets, not a universal law. We recommend re-estimating α periodically (e.g., on a rolling 3-year window).

Sparse maturity grid. The empirical analysis uses 6 Treasury maturities (1Y, 2Y, 5Y, 10Y, 20Y, 30Y). This limits the number of resolvable cosine modes to at most 5 (by Nyquist-type arguments for discrete sampling of continuous basis functions). Higher-resolution data (e.g., the full Treasury curve from GovPX or BrokerTec) would enable testing of higher-order modes ($k \geq 6$) and sharper discrimination between the cosine basis and PCA loadings.

Estimation when the Fejér bound binds. When the unconstrained MLE yields coefficients violating $\sum |A_k| \leq A_0$, the model requires constrained optimization to project onto the feasible set. The paper does not develop the constrained estimation procedure in detail. In practice, the Fejér bound rarely binds for typical yield curve shapes (where $A_0 \gg |A_1|, |A_2|$), but it may bind during extreme inversions. A full treatment of constrained Kalman filtering under the Fejér bound is left for future work.

PCA-cosine basis gap. As discussed in Section 3.2.1, the 99.9% PCA variance figure should not be conflated with the cosine model's fit. The cosine basis approximates the PCA loadings well but is not identical; the approximation quality degrades for higher modes and at boundary maturities.

Nelson–Siegel analogy, not reduction. The spectral model shares the level-slope-curvature interpretation with Nelson–Siegel but uses a different functional basis (cosines vs exponential loadings). Claims of "reducing to" or "subsuming" Nelson–Siegel should be understood as an interpretive analogy, not a mathematical containment.

Independent factor assumption. The model assumes independent OU processes for each mode. In practice, factors exhibit time-varying correlations, particularly during stress periods. Extending the model to correlated factor dynamics (e.g., via a vector OU process) is a natural direction for future work, though it would increase the parameter count.

---

8. Conclusion

The spectral yield curve model achieves three properties simultaneously:

1. **Arbitrage-free by construction** — via the Fejér coefficient bound, not the HJM drift condition. The unweighted bound guarantees yield non-negativity (Lean-verified); the weighted bound provides a sufficient condition for forward-rate positivity. The no-arb check is $O(K)$ per time step.
2. **Strong empirical performance in normal regimes** — PCA explains 99.9% of yield variance with 3 factors; OU dynamics confirmed with $R^2 = 0.98$; the power law $\tau_k \sim k^{-1.3}$ fits with $R^2 = 0.76$ on US Treasury data (2004–2025). The power law is a normal-regime property that breaks down during crises.
3. **Interpretable and parsimonious** — 8 estimable parameters for $K = 3$ (vs 9 for AFNS, ∞ for HJM). The factors are level, slope, and curvature with transparent physical drivers (monetary policy, business cycle, flight-to-quality).

The model shares the Nelson–Siegel factor interpretation while using an orthogonal cosine basis, extends the Diebold–Li framework by adding OU dynamics with a power-law autocorrelation structure $\tau_k = C \cdot k^{-\alpha}$, and replaces the AFNS yield adjustment with a simpler coefficient constraint. The term premium emerges naturally as the risk premium of the most persistent mode.

Acknowledgements

The author used Large Language Models for assistance with exposition. The model, proofs, and empirical analysis are the author's own work. The Lean 4 type checker provides independent verification of the core structural results enumerated in Section 6.1.

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

- Nelson, C.R. and Siegel, A.F (1987). Parsimonious modeling of yield curves. *Nelson, C.R. and Siegel, A.F.*, 60(4). DOI: 10.1086/296409
- Svensson, L.E.O (1994). Estimating and interpreting forward interest rates. NBER Working Paper 4871. *Svensson, L.E.O.*, 1992-1994. DOI: 10.3386/w4871
- Diebold, F.X. and Li, C (2006). Forecasting the term structure of government bond yields. *Diebold, F.X. and Li, C.*, 130(2). DOI: 10.3386/w10048
- Vasicek, O (1977). An equilibrium characterization of the term structure. *Vasicek, O.*, 5(2). DOI: 10.1016/0304-405x(77)90016-2
- Cox, J.C., Ingersoll, J.E., and Ross, S.A (1985). A theory of the term structure of interest rates. *Cox, J.C., Ingersoll, J.E., and Ross, S.A.*, 53(2). DOI: 10.1142/9789812701022_0005
- Heath, D., Jarrow, R., and Morton, A (1992). Bond pricing and the term structure of interest rates. *Heath, D., Jarrow, R., and Morton, A.*, 60(1).
- Duffie, D. and Kan, R (1996). A yield-factor model of interest rates. *Duffie, D. and Kan, R.*, 6(4). DOI: 10.1111/j.1467-9965.1996.tb00123.x
- Christensen, J.H.E., Diebold, F.X., and Rudebusch, G.D (2011). The affine arbitrage-free class of Nelson–Siegel term structure models. *Christensen, J.H.E., Diebold, F.X., and Rudebusch, G.D.*, 164(1). DOI: 10.3386/w13611
- Litterman, R. and Scheinkman, J (1991). Common factors affecting bond returns. *Litterman, R. and Scheinkman, J.*, 1(1). DOI: 10.3905/jfi.1991.692347
- Cont, R. and da Fonseca, J (2002). Dynamics of implied volatility surfaces. *Cont, R. and da Fonseca, J.*, 2(1). DOI: 10.1088/1469-7688/2/1/304
- Filipović, D (1999). A note on the Nelson–Siegel family. *Filipović, D.*, 9(4). DOI: 10.1111/1467-9965.00073
- Björk, T. and Christensen, B.J (1999). Interest rate dynamics and consistent forward rate curves. *Björk, T. and Christensen, B.J.*, 9(1). DOI: 10.1111/1467-9965.00072
- Dai, Q. and Singleton, K.J (2000). Specification analysis of affine term structure models. *Dai, Q. and Singleton, K.J.*, 55(5). DOI: 10.3386/w6128
- Nagy, T. (2026). The Fenton Distribution Solved. *Working paper.*

- Nagy, T. (2026). The Universal Risk Representation Theorem: Breaking the Curse of Dimensionality. *Zenodo*. DOI: 10.5281/zenodo.18910566
- The mathlib Community, “Mathlib: a multi-purpose library for formalized mathematics,” (2020). –2026. https://leanprover-community.github.io/mathlib4_docs/
- Nagy, T. (2026). A Nonlinear Spectral Fundamental Theorem of Asset Pricing. *Working paper.*