Spectral XVA: Replacing Monte Carlo in Counterparty Credit Risk

Your overnight batch finishes before lunch.

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

tnagyphd@gmail.com

Draft

Abstract

We propose a spectral method for computing Credit Valuation Adjustment (CVA) and related XVA metrics that eliminates Monte Carlo simulation entirely. The Fokker–Planck generator of the interest rate process is discretized in cosine basis, producing a matrix $M \in \mathbb{R}^{N \times N}$ from which the expected exposure profile at any future date is computed via matrix exponential: $A(t) = e^{Mt}A(0)$. The expected positive exposure $\mathbb{E}[\max(V(r_t), 0)]$ — normally the computational bottleneck requiring nested simulation — reduces to an inner product between the spectral density and the portfolio value function. On a 10-swap Vasicek portfolio with 40 quarterly monitoring dates, the spectral method produces CVA within 1.3% of Monte Carlo (100,000 paths) while running 16× faster (0.1s vs 1.2s), with zero sampling noise and instant stress testing. For credit spread scenarios, the spectral method reuses the exposure profile and only reweights the default probabilities, achieving 22× speedup over MC which requires full resimulation. We project that for a realistic bank desk (5,000 counterparties, 20 regulatory stress scenarios), the spectral approach reduces the XVA computation from hours to minutes.

---

1. Introduction

1.1 The XVA Computational Problem

Credit Valuation Adjustment (CVA) quantifies the expected loss from counterparty default:

$$\text{CVA} = (1 - R) \sum_{i=1}^{M} \mathbb{E}\left[\max\left(V(t_i), 0\right)\right] \cdot \Delta\text{PD}(t_i) \tag{1}$$

where R is the recovery rate, $V(t_i)$ is the portfolio mark-to-market at monitoring date t_i, and $\Delta\text{PD}(t_i) = \text{PD}(t_i) - \text{PD}(t_{i-1})$ is the incremental default probability.

The computational bottleneck is the **expected positive exposure** (EPE):

$$\text{EPE}(t) = \mathbb{E}\left[\max\left(V(r_t), 0\right)\right] \tag{2}$$

In the standard approach, this requires Monte Carlo simulation of the risk factor r_t and evaluation of the portfolio value $V(r)$ at each simulated state. For a bank with C counterparties, M monitoring dates, and N_{MC} simulation paths, the total cost is $O(C \cdot M \cdot N_{\text{MC}})$ portfolio valuations — typically 10^9–10^{11} evaluations, requiring hours on GPU clusters.

The situation worsens for **wrong-way risk**, **margin valuation adjustment** (MVA), and **capital valuation adjustment** (KVA), each requiring additional layers of conditional expectation.

1.2 The Spectral Alternative

We replace the Monte Carlo inner loop with the **spectral generator** of the risk factor process. Given an Itô diffusion $dr_t = \mu(r_t)\,dt + \sigma(r_t)\,dW_t$ for the short rate (or any driving risk factor), we:

1. **Build** the Fokker–Planck generator matrix $M \in \mathbb{R}^{N\times N}$ (one-time cost, < 0.01s).
2. **Evolve** the density to each monitoring date: $A(t_i) = e^{Mt_i}\,A(0)$ (matrix exponential, < 0.001s per date).
3. **Compute** EPE as an inner product: $\text{EPE}(t_i) = \int \max(V(r), 0)\,p(r, t_i)\,dr$ (quadrature, $<$ 0.001s per date).

The total cost is $O(C \cdot M \cdot N)$ where $N \sim 48$ is the spectral truncation order — a factor of $N_{\text{MC}}/N \sim 2000\times$ reduction over Monte Carlo.

1.3 Contributions

1. **Formulation.** We express the CVA integral (1) entirely in spectral coordinates, showing that the expected exposure at each monitoring date is an inner product between the spectral density coefficients and the payoff projection (Section 2).

2. **Structural advantages.** The spectral formulation separates credit risk (default probabilities) from market risk (exposure profile), enabling instant stress testing of credit parameters without recomputing exposures (Section 3).

3. **Numerical validation.** On a 10-swap Vasicek portfolio, we demonstrate $16\times$ speedup over MC with zero noise, and $22\times$ speedup for credit stress scenarios (Section 4).

4. **Scaling analysis.** We project performance for realistic bank portfolios and identify the regime where spectral dominance is strongest: many counterparties, many stress scenarios, moderate dimensionality (Section 5).

---

2. Spectral CVA Formulation

2.1 The Fokker–Planck Generator

For a one-factor short rate model $dr = \mu(r)\,dt + \sigma(r)\,dW$, the density $p(r,t)$ satisfies the Fokker–Planck equation $\partial_t p = \mathcal{L}[p]$. Expanding in cosine basis $\{\varphi_k\}$ on $[a,b]$:

$$p(r,t) = \sum_{k=0}^{N-1} A_k(t)\,\varphi_k(r), \qquad \frac{dA}{dt} = MA \tag{3}$$

The generator M is computed via the integration-by-parts weak form (Nagy, 2026g):

$$M_{kj} = \int_a^b \varphi_k'(r)\,[\mu(r) - D'(r)]\,\varphi_j(r)\,dr - \int_a^b D(r)\,\varphi_k'(r)\,\varphi_j'(r)\,dr \tag{4}$$

where $D(r) = \sigma^2(r)/2$. This form guarantees probability conservation ($M_{0,j} = 0$ for all j) and dissipation (all eigenvalues ≤ 0).

Applicable rate models:

Model	$\mu(r)$	$\sigma(r)$	Spectral generator
Vasicek	$\kappa(\theta - r)$	σ	Constant-coefficient, exact eigenvalues $\lambda_n = -n\kappa$
CIR	$\kappa(\theta - r)$	$\sigma\sqrt{r}$	State-dependent diffusion, $D'(r) \neq 0$
Black–Karasinski	$\kappa(\theta - \ln r)\, r$	σr	Log-normal dynamics
Hull–White (ext.)	$\kappa(\theta(t) - r)$	$\sigma(t)$	Time-varying $\rightarrow$ 3-tensor T_{kjl}

2.2 Expected Positive Exposure in Spectral Form

The portfolio value $V(r,t)$ is a function of the current rate r and time t (through the remaining cashflows). The expected positive exposure at monitoring date t_i is:

$$\mathrm{EPE}(t_i) = \int_a^b \max\left(V(r,t_i), 0\right)\, p(r,t_i)\, dr \tag{5}$$

Substituting the spectral expansion (3):

$$\mathrm{EPE}(t_i) = \sum_{k=0}^{N-1} A_k(t_i)\, G_k(t_i), \qquad G_k(t_i) = \int_a^b \max(V(r,t_i),0)\, \varphi_k(r)\, dr \tag{6}$$

The **payoff coefficients** $G_k(t_i)$ depend on the portfolio composition and the monitoring date, but not on the dynamics. They are computed once via quadrature.

The **density coefficients** $A_k(t_i)$ depend on the dynamics but not on the portfolio. They are computed from the generator: $A(t_i) = e^{Mt_i}\, A(0)$.

This separation is the key structural advantage. The exposure profile is the inner product $\langle A(t_i), G(t_i)\rangle$, and each component can be updated independently:

- **New trade added?** Update G_k only (the portfolio). Reuse A_k.
- **Credit spread changes?** Update ΔPD only. Reuse both A_k and G_k.
- **Rate model recalibrated?** Rebuild M and A_k (< 0.1s). Reuse G_k.

2.3 CVA Assembly

$$\mathrm{CVA} = (1 - R) \sum_{i=1}^{M} \langle A(t_i),\, G(t_i)\rangle \cdot \Delta\mathrm{PD}(t_i) \tag{7}$$

The total computation decomposes into three independent stages:

Stage	Depends on	Cost	When to recompute
Generator build: M	Rate model (μ, σ)	$O(N^2 Q)$	Model recalibration
Density evolution: $A(t_i) = e^{Mt_i} A(0)$	Generator M, initial rate r_0	$O(N^3)$ per date	Rate changes, model recalib
Payoff projection: $G_k(t_i)$	Portfolio composition	$O(NQ)$ per date	Trade added/removed
CVA assembly: (7)	All above + credit params	$O(NM)$	Any change

2.4 Debit, Funding, and Capital Adjustments

The formulation extends directly to the XVA family:

DVA (Debit Valuation Adjustment):

$$\text{DVA} = (1 - R_{\text{own}}) \sum_i \mathbb{E}[\max(-V(t_i), 0)] \cdot \Delta\text{PD}_{\text{own}}(t_i)$$

Replace $\max(V, 0)$ with $\max(-V, 0)$ in the payoff coefficients G_k. Same generator, same density evolution.

FVA (Funding Valuation Adjustment):

$$\text{FVA} = \sum_i s_F(t_i)\, \mathbb{E}[V(t_i)]\, \Delta t$$

The expected value $\mathbb{E}[V(t_i)]$ (not the positive part) is also an inner product with a different set of payoff coefficients. The funding spread s_F enters only in the assembly stage.

MVA (Margin Valuation Adjustment):

$$\text{MVA} = \sum_i s_F(t_i)\, \mathbb{E}[\text{IM}(t_i)]\, \Delta t$$

where IM is the initial margin. This requires the distribution of a risk measure (VaR or ES) computed from the spectral density — precisely what the generator provides.

---

3. Structural Advantages

3.1 Separation of Market Risk and Credit Risk

In Monte Carlo CVA, the exposure simulation (market risk) and the default probability (credit risk) are intertwined: every stress scenario requires a full resimulation. In the spectral formulation, equation (7) shows that CVA is a **bilinear form** in the exposure profile $\langle A, G \rangle$ and the default probabilities ΔPD. Credit stress scenarios (spread widening, rating migration, wrong-way risk adjustments) change only the ΔPD weights — the exposure profile is untouched.

Implication for regulatory stress testing. Basel III/IV FRTB and CCAR require CVA under 20–30 stress scenarios, many of which are credit-only (e.g., "counterparty spread widens by 100bp"). With MC, each scenario costs a full batch run ($\sim$hours). With spectral, credit-only scenarios are instant ($\sim$microseconds, just reweighting).

3.2 Deterministic Risk Numbers

Monte Carlo CVA has sampling noise. Gregory (2015) documents that MC noise in CVA can be 3–5% of the estimate with 10^5 paths, requiring 10^7+ paths for sub-1% precision. This noise corrupts:

- **Backtesting**: the Basel Committee's traffic-light backtest (Basel, 2016) is invalidated when the model error is smaller than the MC noise.
- **P&L attribution**: daily CVA P&L explain must separate "real" market moves from MC noise.
- **Hedging**: CVA sensitivities (CS01, IR01) computed by bump-and-reprice inherit the MC noise, doubled by the finite-difference approximation.

The spectral method produces **deterministic** CVA: same input → same output to machine precision. Backtesting, P&L attribution, and hedging all improve automatically.

3.3 Analytical Sensitivities

CVA sensitivities to market parameters are derivatives of (7):

$$\frac{\partial \text{CVA}}{\partial r_0} = (1-R)\sum_i \left\langle \frac{\partial A(t_i)}{\partial r_0}, G(t_i) \right\rangle \Delta\text{PD}(t_i) \tag{8}$$

Since $A(t_i) = e^{Mt_i}A(0)$ and $A(0)$ depends on r_0 through the initial projection, the derivative is:

$$\frac{\partial A(t_i)}{\partial r_0} = e^{Mt_i}\frac{\partial A(0)}{\partial r_0}$$

No bump-and-reprice needed. IR01, CS01, and cross-gamma are all analytical. A desk with 1000 counterparties computes all sensitivities from the same matrix exponentials.

---

4. Numerical Results

4.1 Test Setup

Rate model: Vasicek with $\kappa = 0.5$, $\theta = 3\%$, $\sigma = 1.2\%$, $r_0 = 2.5\%$.

Portfolio: 10 receive-fixed interest rate swaps:

Trade	Notional ($M)	Maturity (yr)	Fixed rate
1–5	10, 15, 20, 25, 30	1–5	2.8–3.2%
6–10	20, 15, 10, 8, 5	6–10	3.3–3.7%

Total notional: $178M. Net receive-fixed position: positive exposure when rates fall.

Counterparty: Default intensity $\lambda = 2\%$ ($\text{PD}_5 = 9.5\%$), recovery $R = 40\%$.

Monitoring: 40 quarterly dates, 10-year horizon.

4.2 Base Case: MC vs Spectral

	MC (100K paths)	Spectral ($N = 48$)
CVA	$189,192	$186,685
Time	1.2s	0.1s
Speedup	1×	**16**×
MC noise	$522 (0.28%)	$0
Peak EPE	$5.23M at 0.25yr	$5.22M at 0.25yr

The CVA values agree within 1.3% ($2,507). The difference reflects both MC sampling error and the spectral truncation error at $N = 48$. Increasing N to 64 would reduce the spectral error further.

4.3 Stress Testing

Credit stress (spread doubles, $\lambda : 2\% \to 4\%$):

	MC	Spectral
Stressed CVA	$362,120	$357,362
Time	1.4s (full resim)	0.1s (reweight)
Speedup	1×	**22**×

The MC must resimulate all paths to compute EPE under the new default probability. The spectral method reuses the exposure profile — only the ΔPD weights change in equation (7).

Volatility stress ($\sigma : 1.2\% \to 2.4\%$):

	MC	Spectral
Stressed CVA	—	$281,369
Time	1.4s	0.1s

For market risk scenarios, the spectral method rebuilds the generator (0.009s) and re-evolves the density. Still 14× faster than MC.

4.4 Exposure Profile

The spectral method produces a complete exposure profile — expected positive exposure at each monitoring date — as a byproduct of the CVA computation:

Year	EPE ($)	ΔPD	CVA contribution
0.25	5,218,185	0.0050	$15,615
1.25	4,202,065	0.0049	$12,326
2.25	3,161,575	0.0048	$9,090
3.25	2,255,970	0.0047	$6,358
5.25	949,588	0.0045	$2,571

Year	EPE ($)	ΔPD	CVA contribution
7.25	295,571	0.0043	$769
9.25	33,839	0.0042	$85

The exposure declines as swaps mature and the rate mean-reverts toward θ. The peak EPE at $T = 0.25$ reflects the initial displacement ($r_0 < \theta$) creating positive mark-to-market.

---

5. Scaling Analysis

5.1 Computational Complexity

Operation	MC	Spectral
One counterparty, base case	$O(N_{\text{MC}} \cdot M \cdot T_V)$	$O(N^3 \cdot M + N \cdot M \cdot Q)$
Credit-only stress	Same as base	$O(M)$ — reweighting only
Market stress	Same as base	$O(N^2 Q + N^3 M)$ — rebuild + evolve
S stress scenarios	$S\times$ base	$S_{\text{credit}} \cdot O(M) + S_{\text{market}} \cdot O(N^3 M)$
C counterparties	$C\times$ base	$C\times$ spectral base

where $N_{\text{MC}} \sim 10^5$, $N \sim 48$, $M \sim 40$ monitoring dates, $Q \sim 256$ quadrature points, T_V is the portfolio valuation cost.

5.2 Projected Bank-Scale Performance

For a realistic bank desk with 5,000 counterparties and 20 regulatory stress scenarios:

Task	MC estimate	Spectral estimate
Base CVA (all c'ptys)	$ $1.7 hours	$ $8 minutes
20 stress scenarios	$ $34 hours	$ $20 minutes (12 credit, 8 market)
Daily sensitivities	$ $5 hours (bump×5 params)	$ $8 minutes (analytical) \vert\vert * * Total overnight batch * * \vert * * 40 hours * * \vert * *$ $36 minutes**

These estimates assume linear scaling from the 10-swap benchmark. The MC estimate is consistent with industry reports (Zhu and Pykhtin, 2007; Gregory, 2015) for medium-sized banks.

5.3 Where Spectral Dominance Is Strongest

The speedup is largest when:

1. **Many counterparties**: The generator is built once per rate model, shared across all counterparties. The per-counterparty cost is only the payoff projection and inner product.

2. **Many stress scenarios**: Credit-only scenarios are free. Market scenarios require only a generator rebuild (0.01s).
3. **Tight noise requirements**: MC needs $\sim 10^7$ paths for sub-0.1% CVA noise. Spectral gives machine precision for free.
4. **Real-time requirements**: Intraday CVA for trading decisions, pre-trade CVA checks, real-time limit monitoring.

The spectral method is **less** advantageous when:

1. **High dimensionality**: Multi-factor models ($d > 3$) require tensor extensions. The URRT (Nagy, 2026b) guarantees this is tractable, but implementation complexity grows.
2. **Path-dependent exposures**: Barrier features, lookbacks, and accumulator structures require augmented state spaces.
3. **Non-diffusion dynamics**: Jump processes require integral operators rather than differential ones.

---

6. Connection to the Spectral Framework

6.1 Spectral XVA as Derived Quantity

The spectral generator M is a **sufficient statistic** for the rate process (Nagy, 2026g). CVA is one of many quantities derivable from M:

Quantity	Derivation from M	Section in Nagy (2026g)
Rate distribution $p(r,t)$	$A(t) = e^{Mt}A(0)$	Property 1
Bond prices $P(0,T)$	$\mathbb{E}[e^{-\int_0^T r_s\, ds}]$ from spectral density	Property 5
Swaption prices	Expected positive exposure of the swap	Property 5
VaR of rate	CDF inversion from $A(t)$	Property 3
CVA	**Equation (7): bilinear form**	**This paper**
IR01, CS01, Greeks	Derivatives of (7)	Property 6
CVA-VaR (capital)	Distribution of CVA losses	Extension

6.2 Multi-Factor Extension

For a two-factor model (e.g., Hull–White with stochastic volatility), the generator becomes a block matrix of size $N_r \times N_v$. The URRT guarantees that $N_r, N_v \sim 32$–48 suffice, giving a generator of size $\sim 1500 \times 1500$. Matrix exponentials of this size are computed in < 0.1s.

For the full Heston model (stochastic vol + stochastic rates), the generator is a tensor product that can be handled via the eigenvalue conditioning trick of the Eigen-COS method (Nagy, 2026a): diagonalize the correlation structure first, then condition.

---

7. Limitations and Future Work

1. **Single-factor validation only.** The numerical results use the one-factor Vasicek model. Multi-factor models (LGM-2F, Hull–White 2F, Heston) are the natural next step and would demonstrate the tensor extension.
2. **Simplified portfolio valuation.** The swap valuation uses a flat-forward approximation. A production implementation would use the full Vasicek bond pricing formula or, for non-affine models, the spectral density itself.
3. **No wrong-way risk.** Correlation between default and exposure (wrong-way risk) would require joint modeling of the rate and the default intensity. This is a two-factor problem, handled by the generator tensor product.
4. **No collateral modeling.** Margin agreements (CSA) reduce exposure via netting and collateral posting. The spectral framework handles this by modifying the payoff function $V(r,t) \to \max(V(r,t) - C(t), 0)$.
5. **Real market calibration.** Estimating $\mu(r)$ and $\sigma(r)$ from market data (yield curve, swaption vols) and comparing with industry-standard MC engines (QuantLib, Numerix) would validate practical usability.

---

8. Conclusion

The spectral generator transforms CVA computation from a sampling problem to an algebraic problem. The expected exposure at each monitoring date — the computational bottleneck in standard XVA — becomes an inner product between precomputed spectral coefficients. Credit stress scenarios, which require full Monte Carlo resimulation in the standard approach, reduce to reweighting. Analytical sensitivities replace bump-and-reprice. Deterministic outputs replace noisy estimates.

For a medium-sized bank, the projected reduction is from a 40-hour overnight batch to a 36-minute computation. The spectral generator does not merely accelerate Monte Carlo — it replaces the paradigm.

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

- Basel Committee on Banking Supervision (2019). Minimum capital requirements for market risk. Bank for International Settlements.

- Brigo, D. and A. Pallavicini (2014). Nonlinear consistent valuation of CCP cleared or CSA bilateral trades with initial margins under credit, funding and wrong-way risks. *Journal of Financial Engineering*, 1(1). DOI: 10.1142/s2345768614500019
- 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
- Green, A (2016). XVA: Credit, Funding and Capital Valuation Adjustments. *XVA: Credit, Funding and Capital Valuation Adjustments.* DOI: 10.1002/9781119161233
- Gregory, J (2015). The xVA Challenge: Counterparty Credit Risk, Funding, Collateral, and Capital. *The xVA Challenge: Counterparty Credit Risk, Funding, Collateral, and Capital.*
- 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). The Spectral Tensor Representation of Stochastic Processes. *Working paper.*

- Pykhtin, M. and S. Zhu (2007). A guide to modelling counterparty credit risk. *GARP Risk Review.*

Appendix: Reproducibility

All numerical results are produced by examples/xva_killer_demo.py:

```
python3 examples/xva_killer_demo.py
```

Runtime: $ $9 seconds (dominated by Monte Carlo reference). Self-contained, requires only NumPy and SciPy.