

Signed-Functional Error Bounds and Independent Verification for Operator-Based Option Risk

Daniel Schmitter

Abstract

Numerical error in a signed option-risk response can be bounded more tightly when cancellation is preserved before taking norms. We integrate exact local resolvents, signed Green quadratic forms, whole-spectrum rational time-inversion bounds, and validated residual solves into a compact verifier. A spectral argument bounds continuum time-inversion error independently of the sampled volatility field. The same framework verifies constructive model-ambiguity warnings without trusting the optimizer that proposed them. In a frozen twelve-field synthetic study, four-expiry quote constraints still admit verified material response separation in eight of 24 tested directions. A native recipient subsequently checks supplied positive packets with 42.978 ms warm median latency and 2.391 MiB peak process memory on one observed computation thread. The certificates are conditional on explicit model, observation and floating-point assumptions. They do not certify the true market model, global risk extrema, or narrow uncertainty when no witness is found. The contribution is the validated integration and measured verification interface rather than invention of spectral calculus or interval linear algebra.

1 Introduction

A numerical risk calculation and a risk assertion require different evidence. A point estimate can agree with a reference while its model is weakly identified. Conversely, two materially different risk responses may both be compatible with the same observations. We separate numerical enclosure from model compatibility and use the first to make the second independently checkable.

The core numerical insight is to treat a portfolio and its finite response as signed functionals. Bounding each leg separately discards cancellation before it can help. Exact Green inner products retain that structure and transfer a scalar rational approximation bound to the requested output. A recipient can then recompute a small proposed witness under an independently supplied request.

2 Related work

The spatial construction follows analytic local-volatility and resolvent methods [1]. Rational contour approximations and spectral functional calculus are established [2], as are coefficient-domain inner products [3]. Validated residual techniques are classical floating-point verification [4]. The contribution is their specific combination for signed financial outputs and a tested independent verification interface, not a new general interval theorem or proof-assistant verification of the complete software stack.

3 Architecture and information flow

The verifier consumes a proposition encoded as two model fields, not the proposer’s claimed prices. It recomputes numerical enclosures under a separately trusted request, checks all parameter and quote conditions, and finally checks a signed response separation. Search can be expensive or unreliable without making an accepted witness unreliable, provided the verification assumptions and arithmetic are valid.

The signed-functional construction preserves cancellation before norms are taken. This can tighten an error allowance without accelerating the solver by the same factor. It also separates arithmetic uncertainty from model uncertainty: narrow price intervals can coexist with materially different admissible risk responses. A failed search or rejected packet leaves the question unresolved; only a successful witness establishes the existential ambiguity claim.

Search freely; verify a small assertion



Read left to right above, then right to left below. Arrows show data or computation flow.

Figure 1: Search freely; verify a small assertion. Boxes distinguish supplied information, fitted components, and the quantity evaluated. Arrows show computation or data dependence, not a newly trained deep network.

4 Continuum signed-functional bound

On the interval minus L to L , let K be minus the second derivative plus one quarter with homogeneous Dirichlet endpoints, and M equal two divided by sigma squared. Assume M is bounded above and below by positive constants. The weighted Sturm-Liouville problem has positive eigenvalues and real M -orthonormal eigenfunctions. Its ordinary-coordinate Green resolvent is

$$K\phi_n = \lambda_n M\phi_n, \quad R(s; x, y) = \sum_n \frac{\phi_n(x)\phi_n(y)}{s + \lambda_n}, \quad (1)$$

$$\|\mu\|_{K^{-1}}^2 = \sum_n \frac{\mu(\phi_n)^2}{\lambda_n}. \quad (2)$$

Finite signed point sources and readouts are continuous functionals on the one-dimensional energy space, so these norms are finite. Define T_t by the spectral multiplier $(1 - \exp(-t\lambda))$ divided by λ . Suppose a rational approximation has multiplier error $d(t\lambda)$ divided by λ , with absolute d bounded by δ on the entire nonnegative real axis. Cauchy-Schwarz gives

$$|\nu(T_t\mu) - \nu(\widehat{T}_t\mu)| \leq \delta \|\mu\|_{K^{-1}} \|\nu\|_{K^{-1}}. \quad (3)$$

Indeed, the response error is the sum of $d(t\lambda_n)$ times the source and readout eigen-coefficients divided by λ_n . Bound d uniformly and apply Cauchy-Schwarz to the two weighted coefficient sequences. The norm depends on K , not on M . This is a bounded-domain continuum statement; it does not replace a whole-line truncation allowance or a floating-point evaluation bound.

For point combinations, each norm is a small signed quadratic form in the explicit Green kernel. With p equal to one half, that kernel is

$$G_K(x, y) = \frac{\sinh(p(\min(x, y) + L)) \sinh(p(L - \max(x, y)))}{p \sinh(2pL)}, \quad (4)$$

$$\left\| \sum_i q_i \delta_{x_i} \right\|_{K^{-1}}^2 = \sum_{i,j} q_i q_j G_K(x_i, x_j). \quad (5)$$

The spread source uses weighted opposite-sign atoms at log strikes minus 0.05 and 0.05. A fixed-coordinate finite spot response uses a weighted atom at the bumped spot minus the original atom. Computing these

signed norms before combining them retains cancellation. The reported initial example tightens the transferred bound by about 172 times relative to adding individual leg bounds; this is bound tightness, not execution speed.

5 Time inversion and numerical enclosure

The scalar certificate treats the serialized binary64 contour coefficients as exact inputs. Interval Taylor expansions cover finite spectral intervals; an inverse-variable expansion encloses the infinite tail. Remainders are included, rather than replacing the supremum with sampled values. The application’s resulting scalar bound is approximately 4.249×10^{-12} .

For each complex tridiagonal solve $Au = b$, an approximate solution and inverse C propose a correction. If an outward bound η for the infinity norm of $I - CA$ is below one, the Neumann-series argument gives

$$\|u - \tilde{u}\|_\infty \leq \frac{\|C(b - A\tilde{u})\|_\infty}{1 - \eta}, \quad (6)$$

$$\eta \geq \|I - CA\|_\infty < 1. \quad (7)$$

All products and residuals on the right must be enclosed. The exact cell formulas remove spatial discretization error for the declared layered model, but elementary-function rounding, assembly, rational time error, readout arithmetic, and serialization remain. Native code assumes IEEE binary64, nearest rounding, gradual underflow, and strict compiler settings. Independent high-precision interval references provide a separate check, not a formal proof of the implementation.

A global volatility cap and declared coefficient extension support a separate boundary-hit/heat allowance when a whole-line claim is required. The cap is an assumption outside the observed domain too. For changing slabs, a common energy-coordinate contraction argument controls products of semigroups; the implemented two-slab extension needs confluent cross-products. It passes a fresh 32-case accuracy panel but misses its 25 ms timing target at 28.26 ms. No arbitrary-depth cost claim follows.

6 Observation and witness verification

The synthetic class has fifteen log-volatility cells, volatility between 0.08 and 0.60, and adjacent log contrasts at most 0.70. The normalized martingale uses the declared $\exp(\text{minus } 0.01 \text{ times maturity})$ price discount. Four expiries are 0.01, 0.05, 0.2 and 0.5; the risk target is the spread’s finite response at 0.05 to minus or plus 0.5% spot moves with the coefficient field fixed.

Each expiry has 21 quote coordinates and two spread-leg quotes. Gaussian synthetic noise uses a recorded strike-dependent scale with minimum 5×10^{-5} . Eight discrepancy blocks have fixed binary64 chi-square thresholds at probability one minus 0.01/8. The nominal family-wise interpretation assumes that noise model; the numerical verifier checks serialized limits, not the correctness of a market-noise assumption or the statistical quantile algorithm.

A witness contains two coefficient fields and one direction. Acceptance requires both complete enclosed quote discrepancies to lie within every applicable limit, their parameters and numerical widths to qualify, and their enclosed response separation to exceed 2×10^{-4} . This proves existence of two separated compatible model responses. It does not prove the global minimum or maximum. Failure to find or verify a pair leaves the direction unresolved.

7 Experimental methods

Twelve fresh synthetic fields consist of two constants, two alternating controls, and eight reflected log-volatility walks. Generation and observation seeds are 2026091719 and 2026091720. Both single-expiry and shared-four-expiry searches use the same two feasibility starts and eight signed-response attempts, with fixed SLSQP iteration and evaluation budgets. Every selected endpoint is independently checked at all maturities; rejected endpoints are not replaced. Optimizer failure is retained separately from whether a visited feasible point supplies a valid witness.

The subsequent deployment experiment reuses 42 exposed positive/negative packet cases and their reference enclosures. A packet binds its fields and direction to the digest of an independently trusted request; it carries no prices to be trusted. The native recipient recomputes eight field/maturity combinations, parameter conditions, quote blocks, allowances, and risk separation. Three timed warm repetitions, one retained warm-up, one cold process per case, and stale-identity/recovery tests are recorded. Positive-case timings are reported separately from cheaper rejection paths.

8 Results

Table 1: Frozen synthetic comparison. No significance claim is made from twelve model fields.

Outcome	Single expiry	Shared four expiries
Material directions / 24	18	8
Admitted selected endpoints	48 / 48	47 / 47
Failed optimizer attempts / 120	39	52
Median search + verification (s)	0.392	1.371




Figure 2: Verified material directions in the single- and four-expiry synthetic studies. Missing witnesses remain unresolved; these counts are not upper bounds on model ambiguity.

All twelve generating fields are admitted. None of the single-expiry selected fields satisfies the full four-expiry contract. The eight joint warnings occur in four random fields, in both spot directions, with guaranteed separations approximately 0.000238-0.000524. The other sixteen directions are unresolved, not certified narrow. Thus additional expiries reduce the found ambiguity but do not eliminate it in this bounded search. The unchanged single-slab numerical engine contains independent references in all 320 fresh field/maturity cases, with maximum price width about 6.18×10^{-10} . In the later single-thread deployment, all 210 warm-up/warm/cold packet classifications and their applicable reference checks pass. Positive warm median is 42.978 ms, p95 44.956 ms, and cold p95 48.673 ms. Peak native memory is 2.391 MiB and the executable is 123,096 bytes. These are complete request measurements on the shared workstation, not search cost or weak-device results.

9 Discussion

A digest prevents a mismatched packet from being applied to a different trusted request; it does not authenticate an exchange or make its data reliable. Likewise, numerical certificates do not establish that a

local diffusion explains a market. The recipient is a bounded research prototype, not a production security assessment.

A full-class upper risk certificate remains absent. Parameter-box bounds become impractical over broad ranges, and a relaxed positive-operator class has provably wide completions in the tested diagnostic. A constructive warning is useful precisely because it is a smaller assertion that can be independently checked. It should never be turned into reassurance when no warning is found.

10 Application boundary and research implication

This architecture is relevant beyond pricing wherever a complex optimizer proposes a small checkable object. Its scope here is constructive warning, not a certificate that risk is narrow over the entire class. Request identity protects against accidental mismatch; it is not authentication of market data or a security audit of a production service.

11 Conclusion

Exact signed inner products and spectral bounds turn a supplied pair of candidate models into a numerically checkable ambiguity warning at small native runtime cost. The integrated result separates numerical precision, observational compatibility, and model uncertainty. Its practical value remains conditional on an adequate observation/model interface and the decision that consumes the warning.

References

- [1] A. Itkin and A. Lipton. Filling the gaps smoothly. *Journal of Computational Science*, 2018; author preprint, 2016. <https://arxiv.org/abs/1608.05145>
- [2] L. N. Trefethen, J. A. C. Weideman, and T. Schmelzer. Talbot Quadratures and Rational Approximations. *BIT* 46, 653-670, 2006. <https://doi.org/10.1007/s10543-006-0077-9>
- [3] A. Badoual, D. Schmitter, and M. Unser. An Inner-Product Calculus for Periodic Functions and Curves. *IEEE Signal Processing Letters* 23(6), 878-882, 2016. <https://doi.org/10.1109/LSP.2016.2555139>
- [4] S. M. Rump. Verification Methods: Rigorous Results Using Floating-Point Arithmetic. *Acta Numerica* 19, 287-449, 2010. <https://doi.org/10.1017/S096249291000005X>