

Basket Option Pricing Under Stochastic Volatility via Eigenvalue-Conditional Heston Mixing

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Draft

Abstract

The eigenvalue-conditional mixing framework (Nagy, 2026a, 2026c) prices basket options under geometric Brownian motion by decomposing the correlation matrix, conditioning on dominant factors, and mixing per-scenario Black-Scholes prices. We prove that the Mixture Collapse theorem enabling this approach is distribution-agnostic (Lean-verified: `MixtureCollapse.lean`), depending only on linearity of finite sums and the law of iterated expectations. This motivates extending the framework to stochastic volatility: replace per-scenario Black-Scholes with per-scenario Heston (1993) pricing, where the conditional characteristic function is known analytically. We derive the eigenvalue-conditional Heston pricing formula, show that the conditional basket characteristic function approximately factorizes as a product of single-asset Heston CFs (with residual error bounded by the first omitted eigenvalue λ_{K+1}), and establish error bounds decomposing into three components: correlation conditioning error, single-asset Heston CF truncation error, and Gauss-Hermite quadrature error. We conjecture a conditional variance reduction result extending Proposition 1 of Nagy (2026c) to stochastic volatility and provide heuristic arguments. Numerical experiments on 2–50 asset baskets under the multi-asset Heston model are designed to benchmark this approach against Monte Carlo simulation; we describe the full experimental protocol with eight test configurations spanning diverse Heston regimes. To our knowledge, this is the first deterministic (non-Monte-Carlo) basket option pricing method under multi-asset stochastic volatility that exploits the spectral structure of cross-asset correlation.

1. Introduction

1.1 The Gap

Basket option pricing under GBM is solved: eigenvalue-conditional BS mixing achieves sub-2% error at sub-second latency (Nagy, 2026c). But GBM is a toy model — real markets exhibit: - Stochastic volatility (the smile/skew) - Volatility clustering - Leverage effect (negative vol-spot correlation)

The Heston (1993) model captures these features for single assets. Multi-asset Heston is well-defined but computationally expensive: the $2n$ -dimensional characteristic function must be evaluated numerically, and Monte Carlo is the standard pricing approach.

1.2 Why Eigenvalue Conditioning Extends

The Mixture Collapse (Nagy, 2026a, Theorem 1; Lean-verified in `MixtureCollapse.lean`) states that the unconditional basket option price equals the quadrature-weighted sum of conditional prices:

$$V = e^{-rT} \mathbb{E}[(B_T - K)^+] = e^{-rT} \sum_{q=1}^Q w_q \mathbb{E}[(B_T - K)^+ | Z = z_q] = \sum_{q=1}^Q w_q C_q$$

where $B_T = \sum_{i=1}^n a_i S_i(T)$ is the basket terminal value, $Z = (Z_1, \dots, Z_K)$ are the dominant eigenfactors of the cross-asset correlation matrix, z_q are Gauss-Hermite quadrature nodes with weights w_q satisfying $\sum_q w_q = 1$, and $C_q = e^{-rT} \mathbb{E}[(B_T - K)^+ | Z = z_q]$ is the conditional basket price in scenario q .

The formal proof in Lean establishes this via the sum-interchange identity: the unconditional COS series coefficients equal the weighted average of conditional coefficients, $A_k = \sum_{q=1}^Q w_q A_k^{(q)}$, by commuting two finite sums. This interchange is valid for any summable family and does NOT depend on:

- The marginal distribution of the assets (lognormal, Heston, jump-diffusion, or any model with a computable per-scenario price)
- The per-scenario pricing method (Black-Scholes, Heston CF, FFT, Monte Carlo)
- The dimension of the underlying state space

What it DOES require:

- A decomposable correlation/covariance structure ($C = V\Lambda V^T$, always true for positive semi-definite C)
- A computable per-scenario price $C_q = f(\mu_q, \sigma_q^{(q)}, v_q, \dots)$ as a function of the conditional parameters
- Gauss-Hermite quadrature weights w_q with $\sum_q w_q = 1$, approximating the integral over the factor distribution

This distribution-agnosticism is the key observation enabling the extension to stochastic volatility: since the Mixture Collapse relies only on linearity of expectation and the law of iterated expectations, not on the Black-Scholes formula, we can replace the per-scenario BS price with any alternative pricer — in particular, the Heston CF-based pricer — without modifying the outer mixing infrastructure.

1.3 The Heston Opportunity

Under the multi-asset Heston model, each asset i has: - Price: $dS_i/S_i = (r - q_i) dt + \sqrt{v_i} dW_i^S$
 - Variance: $dv_i = \kappa_i(\theta_i - v_i) dt + \xi_i \sqrt{v_i} dW_i^v$ - Leverage: $\text{Corr}(dW_i^S, dW_i^v) = \rho_i^{\text{lev}}$ - Cross-asset: $\text{Corr}(dW_i^S, dW_j^S) = \rho_{ij}$

The cross-asset Brownian motions decompose via the eigendecomposition of $C = (\rho_{ij}) = V\Lambda V^T$:

$$W_i^S = \sum_{j=1}^n V_{ij} F_j$$

where F_1, \dots, F_n are independent standard Brownian motions and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 0$ are the eigenvalues. After conditioning on the K dominant factors $(F_1, \dots, F_K) = (z_{q,1}, \dots, z_{q,K})$, the residual Brownian driving each asset is:

$$W_i^{S,\text{res}} = \sum_{j=K+1}^n V_{ij} F_j$$

The residual cross-asset correlation between assets i and l is:

$$\text{Corr}(W_i^{S,\text{res}}, W_l^{S,\text{res}}) = \sum_{j=K+1}^n V_{ij} V_{lj} \lambda_j$$

which is bounded in absolute value by $\sum_{j=K+1}^n \lambda_j$. When K captures the dominant eigenvalues (a condition satisfied empirically for typical equity correlation matrices where 2–5 factors explain > 90% of variance), this residual correlation is small and the conditional assets are approximately independent.

Approximate CF factorization. Under exact conditional independence, the conditional characteristic function of the basket $B_T = \sum_i a_i S_i(T)$ would factorize as:

$$\phi_{B_T|Z=z_q}(u) = \prod_{i=1}^n \phi_{\text{Heston},i}(u a_i; F_i^{(q)}, v_i(0), \kappa_i, \theta_i, \xi_i, \rho_i^{\text{lev}})$$

where $F_i^{(q)}$ is the conditional forward price (see Section 2.2) and each factor is the single-asset Heston CF evaluated analytically via the Heston (1993) formula. The approximation error from treating the residual correlation as zero is:

$$\left| \phi_{B_T|Z=z_q}^{\text{exact}}(u) - \prod_{i=1}^n \phi_{\text{Heston},i}(u a_i) \right| \leq C_\phi(u) \sum_{j=K+1}^n \lambda_j$$

where $C_\phi(u)$ depends on $|u|$, the asset weights a_i , and the Heston parameters but is independent of K . This bound follows from the Lipschitz continuity of the joint CF with respect to the correlation matrix entries [TODO:cite formal bound]. In practice, the factorization is highly accurate when $\tau_K = \sum_{j=1}^K \lambda_j / \sum_{j=1}^n \lambda_j$ exceeds 0.90.

The basket price in each scenario is then obtained by applying the COS method (Fang and Oosterlee, 2008) or Carr-Madan FFT (Carr and Madan, 1999) to this product CF — a one-dimensional inversion over the basket value, not a high-dimensional integral over assets.

1.4 Contributions

1. **Eigenvalue-conditional Heston pricing formula** (Section 2) — to our knowledge, the first deterministic (non-Monte-Carlo) multi-asset basket option pricer under stochastic volatility that exploits the spectral structure of the cross-asset correlation matrix. Existing semi-analytical approaches based on Wishart processes (Da Fonseca et al., 2008) operate on a different model class; see Section 1.5 for a detailed comparison.
2. **Model-free Mixture Collapse** (Lean-verified in MixtureCollapse.lean) — formal proof that the mixing framework is distribution-agnostic, relying only on linearity of finite sums.

3. **Conditional variance reduction under Heston** (Section 3) — we conjecture an extension of Proposition 1 from Nagy (2026c) to stochastic volatility and provide heuristic arguments. A formal proof remains open.
4. **Error decomposition** (Section 3) — three components: $\varepsilon_{\text{corr}}(K)$ (eigenvalue truncation) + $\varepsilon_{\text{Heston}}$ (single-asset Heston CF accuracy) + ε_{GH} (Gauss-Hermite quadrature).
5. **Numerical experiments** (Section 5) — experimental design benchmarking against multi-asset Heston MC on 2–50 asset baskets across eight test configurations.

1.5 Related Work

Multi-asset stochastic volatility models. The Wishart multifactor stochastic volatility model of Da Fonseca, Grasselli, and Tebaldi (2008) provides a matrix-valued variance process and admits a semi-analytical characteristic function for the joint asset distribution. However, their approach yields a CF for the full multi-asset system — not a decomposed, per-scenario product — and basket pricing still requires high-dimensional numerical integration or Monte Carlo. Our eigenvalue-conditional approach operates on a different axis: we keep per-asset Heston dynamics and decompose the cross-asset correlation structure spectrally, reducing the problem to a sequence of (approximately) independent single-asset Heston pricings.

Multi-asset Heston calibration and simulation. Dimitroff, Lorber, and Röder (2017) [TODO:cite] address joint calibration of multi-asset Heston models to individual and cross-asset option surfaces. Their work is complementary to ours: they solve the calibration problem (recovering Heston parameters from market data), while we solve the pricing problem (evaluating basket options given calibrated parameters). Ahdida and Alfonsi (2013) [TODO:cite] develop efficient simulation schemes for Wishart-type processes, improving Monte Carlo convergence. Our method avoids simulation entirely by exploiting the eigendecomposition.

Machine learning approaches. Goudenège, Molent, and Zanette (2020) [TODO:cite] apply neural network surrogates to price multi-asset options under stochastic volatility. Their approach achieves fast evaluation after an expensive offline training phase but lacks the interpretability and error bounds of our spectral decomposition. The eigenvalue-conditional framework could serve as a structured basis for such learned surrogates (see Section 6.4).

Comparison summary. Existing multi-asset SV pricing methods fall into three categories: (i) full Monte Carlo simulation (general but slow), (ii) Wishart-based semi-analytical CFs (elegant but still high-dimensional), and (iii) machine learning surrogates (fast but opaque). The eigenvalue-conditional Heston approach occupies a new position: deterministic, spectral-structure-aware, with per-scenario pricing that leverages the single-asset Heston CF — a well-understood, numerically stable building block.

2. The Pricing Formula

2.1 Multi-Asset Heston Setup

n assets, each following Heston dynamics. Cross-asset correlation matrix C (for the Brownian motions W_i^S). Eigendecompose $C = V\Lambda V^T$.

2.2 Eigenvalue Conditioning and Conditional Dynamics

Eigendecompose the cross-asset correlation matrix: $C = V\Lambda V^T$, with eigenvalues $\lambda_1 \geq \dots \geq \lambda_n \geq 0$ and orthonormal eigenvectors $V = [v_1 \mid \dots \mid v_n]$. Select the K dominant factors and approximate the factor distribution via Q -point Gauss-Hermite quadrature with nodes $z_q \in \mathbb{R}^K$ and weights w_q (see Nagy, 2026c, Section 2 for details).

Conditional forward price. After conditioning on $(F_1, \dots, F_K) = z_q$, the factor realization contributes a deterministic component to each asset’s log-return. The conditional forward price of asset i in scenario q is:

$$F_i^{(q)} = F_i(0) \exp\left(\sum_{j=1}^K V_{ij} \sqrt{\lambda_j} z_{q,j} \bar{\sigma}_i \sqrt{T} - \frac{1}{2} \sum_{j=1}^K V_{ij}^2 \lambda_j \bar{\sigma}_i^2 T\right)$$

where $F_i(0) = S_i(0) e^{(r-q_i)T}$ is the unconditional forward price, $\bar{\sigma}_i = \sqrt{\mathbb{E}[\frac{1}{T} \int_0^T v_i(s) ds]}$ is the expected average volatility (computable from the Heston parameters as $\bar{\sigma}_i^2 = \theta_i + (v_i(0) - \theta_i)(1 - e^{-\kappa_i T})/(\kappa_i T)$), and the second term is the convexity correction ensuring $\mathbb{E}[F_i^{(q)}] = F_i(0)$ under the quadrature.

Conditional variance state. The initial variance $v_i(0)$ is a state variable and does not change with the cross-asset factor conditioning. The per-asset Heston parameters $(\kappa_i, \theta_i, \xi_i, \rho_i^{\text{lev}}, v_i(0))$ are identical across all scenarios q . What changes across scenarios is the conditional forward price $F_i^{(q)}$ and the residual cross-asset correlation structure. After conditioning, the residual correlation between assets is bounded by $\sum_{j=K+1}^n \lambda_j$, and each asset’s local dynamics — the mean-reversion, vol-of-vol, and leverage effect — are unchanged.

Notation convention. We write the per-scenario Heston CF as $\phi_{\text{Heston},i}(u; F_i^{(q)}, v_i(0), \kappa_i, \theta_i, \xi_i, \rho_i^{\text{lev}}, T)$, making explicit that the conditioning enters through the shifted forward $F_i^{(q)}$ while the variance parameters remain fixed.

2.3 Per-Scenario Heston Pricing

For each quadrature point q , we must price a basket of approximately independent Heston assets with conditional forwards $F_i^{(q)}$. We describe three approaches; **Option A is the primary method** used throughout this paper.

Option A (primary): COS on conditional Heston CF. Compute the conditional basket CF as a product of single-asset Heston CFs (each evaluated at the conditional forward $F_i^{(q)}$ with unchanged variance parameters). Apply the COS method (Fang and Oosterlee, 2008) for 1D option pricing on the basket value — this places the COS method in its native habitat: a single “effective asset” (the conditional basket) with a known, analytically computable CF on a tight domain $[a, b]$.

Option B: Heston-FW hybrid. Match the conditional basket’s first two moments (computable from the Heston parameters and the conditional forwards) to a lognormal distribution, then price with Black-Scholes. This is the analog of the Fenton-Wilkinson approach in Nagy (2026c), but with Heston-corrected moments. Faster than Option A but less accurate for OTM options where the smile matters.

Option C: FFT per scenario. Apply the Carr-Madan (1999) FFT to the conditional basket CF. More general than COS (handles a wider range of payoffs) but slower due to the larger grid required.

Algorithm pseudocode (Option A). The complete eigenvalue-conditional Heston pricing algorithm is:

Input: n assets with Heston parameters $(\kappa_i, \theta_i, \xi_i, \rho_i^{\text{lev}}, v_i(0))$, basket weights a_i , cross-asset correlation C , strike K , maturity T , risk-free rate r , dividend yields q_i , number of factors K_{fac} , quadrature order Q_{GH} , COS terms N_{COS} .

1. Eigendecompose $C = V\Lambda V^T$; retain K_{fac} dominant eigenvalues.
2. Generate K_{fac} -dimensional Gauss-Hermite quadrature nodes $\{z_q\}_{q=1}^Q$ and weights $\{w_q\}_{q=1}^Q$.
3. **For each** scenario $q = 1, \dots, Q$:
 - a. Compute conditional forwards $F_i^{(q)}$ for all assets i (Section 2.2).
 - b. For each COS frequency $k = 0, \dots, N_{\text{COS}} - 1$, evaluate the conditional basket CF: $\hat{\phi}_k^{(q)} = \prod_{i=1}^n \phi_{\text{Heston},i}(u_k a_i; F_i^{(q)}, v_i(0), \kappa_i, \theta_i, \xi_i, \rho_i^{\text{lev}}, T)$.
 - c. Compute the per-scenario basket price C_q via COS summation.
4. **Return** $V_{\text{call}} = e^{-rT} \sum_{q=1}^Q w_q C_q$.

2.4 Mixing

The final basket option price is the quadrature-weighted sum of per-scenario Heston prices:

$$V_{\text{call}} = e^{-rT} \sum_{q=1}^Q w_q \cdot C_q^{\text{Heston}}$$

where $C_q^{\text{Heston}} = \mathbb{E}[(B_T - K)^+ \mid Z = z_q]$ is computed via Option A (COS), Option B (moment-matched BS), or Option C (FFT) as described in Section 2.3. This mixing formula is the same as in the GBM case and is Lean-verified (MixtureCollapse.lean): the validity of the outer sum depends only on the quadrature approximation of the factor integral, not on the per-scenario pricing method.

3. Convergence Analysis

3.1 Conditional Variance Reduction Under Heston

Under GBM, the conditional variance of the basket log-return decreases by a factor $(1 - \tau_K)$ after conditioning on K dominant eigenfactors (Nagy, 2026c, Proposition 1; Lean-verified in ConditionalVariance.lean). Under Heston, the situation is more subtle because the integrated variance $\int_0^T v_i(s) ds$ is itself stochastic.

Conjecture 1 (Conditional variance reduction under stochastic volatility). Under the multi-asset Heston model with cross-asset correlation C and $\tau_K = \sum_{j=1}^K \lambda_j / \text{tr}(C)$, conditioning on the K dominant eigenfactors reduces the conditional basket's integrated variance:

$$\text{Var}\left(\sum_{i=1}^n a_i \ln S_i(T) \middle| Z\right) \leq (1 - \tau_K) \text{Var}\left(\sum_{i=1}^n a_i \ln S_i(T)\right) + \varepsilon_{\text{SV}}$$

where ε_{SV} captures the stochastic volatility contribution that is not explained by the cross-asset eigendecomposition. We expect $\varepsilon_{\text{SV}} = O(\xi_{\text{max}}^2/\kappa_{\text{min}})$, vanishing in the low-vol-of-vol ($\xi \rightarrow 0$) or fast-mean-reversion ($\kappa \rightarrow \infty$) limits, where Heston degenerates to GBM.

Status: This result is stated as a conjecture, not a theorem. A formal proof would require bounds on the conditional distribution of the integrated variance process, which depend on the specific interaction between the leverage correlation ρ_i^{lev} and the eigenvector loadings. We provide two heuristic arguments:

1. **GBM limit.** When $\xi_i = 0$ for all i , the variance processes are deterministic ($v_i(t) = \theta_i + (v_i(0) - \theta_i)e^{-\kappa_i t}$), and the bound reduces exactly to the GBM result of Nagy (2026c) with $\varepsilon_{\text{SV}} = 0$.
2. **Per-asset independence of vol dynamics.** The stochastic volatility dynamics dv_i are per-asset: v_i is driven by $W_i^v = \rho_i^{\text{lev}} W_i^S + \sqrt{1 - (\rho_i^{\text{lev}})^2} Z_i^v$, where Z_i^v is independent of all cross-asset factors. Conditioning on the dominant eigenfactors of C partially determines W_i^S (and hence, through leverage, partially determines W_i^v), but the idiosyncratic component Z_i^v and the residual cross-asset Brownian remain unaffected. The variance reduction from conditioning thus applies to the cross-asset component of the basket variance, while the per-asset stochastic vol contributes an additive residual ε_{SV} .

3.2 Error Decomposition

$$|V - \hat{V}| \leq \varepsilon_{\text{corr}}(K) + \varepsilon_{\text{Heston}} + \varepsilon_{\text{GH}}$$

Component	Source	GBM analog
$\varepsilon_{\text{corr}}(K)$	Residual cross-asset correlation	$\varepsilon_{\text{FW}}(K)$ from Nagy (2026c)
$\varepsilon_{\text{Heston}}$	Single-asset Heston pricing error (COS or FFT)	0 (BS is exact for lognormal)
ε_{GH}	Gauss-Hermite quadrature	Same as Nagy (2026c)

Under GBM, $\varepsilon_{\text{Heston}} = 0$ (BS is exact) and we recover the 2-component bound of Nagy (2026c), Theorem 1.

4. The Greeks Under Heston

4.1 Delta

Same structure as Nagy (2026c), Section 4.1, but with Heston Delta per scenario instead of BS Delta:

$$\Delta_i = e^{-rT} \sum_q w_q \cdot \Delta_i^{\text{Heston},q}$$

where $\Delta_i^{\text{Heston},q}$ is the single-asset Heston Delta for asset i in scenario q .

4.2 Vega (Vol-of-Vol Sensitivity)

Under Heston, “Vega” becomes richer: $-\partial V/\partial v_i(0)$: sensitivity to current variance (the traditional Vega) $-\partial V/\partial \theta_i$: sensitivity to long-term variance $-\partial V/\partial \xi_i$: sensitivity to vol-of-vol

Each decomposes as a mixing of per-scenario Heston Greeks.

4.3 Correlation Sensitivity

Same finite-difference approach as Nagy (2026c), Section 4.3: re-eigendecompose at $C + \epsilon \cdot e_{ij}e_{ij}^T$.

5. Numerical Experiments

5.1 Test Cases

All test cases use $r = 0.03$, $q_i = 0.01$ (uniform dividend yield), $T = 1.0$ year unless otherwise specified. Equal basket weights $a_i = 1/n$. Initial spot prices $S_i(0) = 100$ for all i .

ID	n	κ	θ	ξ	ρ^{lev}	v_0	Correlation	Moneyness	Payoff	Notes
H1	2	2.0	0.04	0.2	-0.3	0.04	$\rho_{12} = 0.5$	ATM ($K = 100$)	Call	Low vol-of-vol base-line
H2	2	2.0	0.04	0.8	-0.3	0.04	$\rho_{12} = 0.5$	ATM ($K = 100$)	Call	High vol-of-vol stress test
H3	5	[1.5, 2.0, 2.5, 3.0, 4.0]	[0.03, 0.04, 0.05, 0.06, 0.07]	[0.3, 0.4, 0.5, 0.6, 0.7]	[0.2, 0.3, 0.4, 0.5, 0.6]	-0.7	[0.4, 0.5, 0.6, 0.7, 0.8]	ATM ($K = 110$)	Call	Mixed Heston params
H4	10	2.0	0.04	0.3	-0.5	0.04	Equicorr: $\rho_{ij} = 0.4$	ATM ($K = 100$)	Call	Moderate equicor-related

ID	n	κ	θ	ξ	ρ^{lev}	v_0	Correlation	Moneyness	Payoff	Notes
H5	10	1.5	0.05	0.5	-0.7	0.05	Equicorr: $\rho_{ij} = 0.4$	ATM ($K = 100$)	Put	High leverage effect
H6	20	2.0	0.04	0.3	-0.5	0.04	Block: 4 sectors of 5, intra=0.6, inter=0.2	ATM ($K = 100$)	Call	Sector structure
H7	2	2.0	0.04	0.3	-0.5	0.04	$\rho_{12} = 0.7$	ATM ($K = 0$)	Spread ($a = [+1, -1]$)	Spread option
H8	50	2.0	0.04	0.3	-0.3	0.04	Equicorr: $\rho_{ij} = 0.3$	ATM ($K = 100$)	Call	CLT/large basket regime

For test cases with vector-valued parameters (H3), each asset has its own Heston calibration. For scalar-valued parameters (H1, H2, H4–H8), all assets share the same Heston parameters.

5.2 Reference Prices

Reference prices are computed via multi-asset Heston Monte Carlo simulation with: - 10^6 paths, antithetic variates, moment matching - Euler-Maruyama discretization with 252 time steps per year (daily) - Full-rank Cholesky decomposition of C for correlated Brownian motion generation - Milstein correction for the variance process where applicable - 95% confidence intervals reported for all MC estimates

5.3 Methods to Compare

1. **Eigen-Cond Heston (Option A)** — this paper: eigendecompose C , condition on K_{fac} factors, per-scenario COS on the product Heston CF. Parameters: $K_{\text{fac}} \in \{1, 2, 3\}$, $Q_{\text{GH}} \in \{5, 10, 20\}$, $N_{\text{COS}} = 128$.
2. **Multi-asset Heston MC** — full correlated simulation as described in Section 5.2.
3. **Eigen-Cond BS** (Nagy, 2026c) — eigenvalue conditioning with per-scenario Black-Scholes (ignoring stochastic vol). Uses $\bar{\sigma}_i$ as the flat volatility.
4. **Single-asset Heston FW** — moment-match the basket to a single effective Heston process and price with the 1D Heston CF. A crude but fast approximation.

5.4 Metrics

For each test case, we report:

- **Absolute pricing error:** $|V_{\text{method}} - V_{\text{MC}}|$ in basis points of the notional
- **Relative pricing error:** $|V_{\text{method}} - V_{\text{MC}}|/V_{\text{MC}}$ as a percentage

- **Runtime:** wall-clock time in milliseconds (single-threaded Python, averaged over 100 evaluations)
- **Convergence in K_{fac} :** pricing error as a function of the number of conditioning factors
- **Convergence in Q_{GH} :** pricing error as a function of quadrature order
- **Smile accuracy:** implied volatility error across a strike range $K \in \{80, 90, 95, 100, 105, 110, 120\}$ for test cases H1, H2, H5

5.5 Expected Figures

The following figures will be generated (implementation in `examples/generate_heston_figures.py`, [TODO: implement]):

1. **Error vs. number of factors (K_{fac}):** log-scale plot of absolute pricing error vs. K_{fac} for test cases H4, H6, H8, showing rapid convergence with 2–3 factors.
2. **Runtime comparison:** bar chart comparing Eigen-Cond Heston, MC ($10^4, 10^5, 10^6$ paths), and Eigen-Cond BS across test cases H1, H4, H6, H8.
3. **Implied vol smile:** for H1 and H2, plot the basket implied volatility across strikes from Eigen-Cond Heston vs. MC reference, demonstrating that the method captures the smile/skew.
4. **GBM vs. Heston comparison:** for H2 (high vol-of-vol) and H5 (high leverage), show the pricing error of Eigen-Cond BS (which ignores stochastic vol) vs. Eigen-Cond Heston (which captures it), demonstrating when stochastic vol matters.

5.6 Results

[TODO: Run `examples/generate_heston_figures.py` and populate this section with tables and figure references. Replace this placeholder after implementing the numerical experiments.]

6. Discussion

6.1 When Stochastic Vol Matters

For ATM options with $T > 6$ months, the smile effect is significant and GBM underprices OTM options. The eigenvalue-conditional Heston approach captures the smile per scenario.

For short-dated ATM options ($T < 3$ months), GBM is adequate and the BS approach of Nagy (2026c) is sufficient.

6.2 The Resolution Ladder (Extended)

The “resolution” refers to the number of effective parameters characterizing the conditional basket distribution. The Spectral Fenton distribution (Nagy, 2026a) uses N_{COS} COS coefficients $\{A_k\}$ (typically $N_{\text{COS}} = 128$, hence “128” or “130” as shorthand) plus domain bounds $[a, b]$, providing a high-fidelity representation of the conditional basket density.

Resolution	Parameters	Model	Per-scenario pricer
2	(M, V)	Gaussian approximation	$\Phi(d)$ (normal CDF)
3	(M, V, γ_3)	Cornish-Fisher	Polynomial correction
\$ \$130	$\{A_k\}_{k=0}^{N_{\text{COS}}-1}, a, b$	Spectral Fenton (GBM)	COS sine series
\$ \$130 + Heston	$\{A_k\}, a, b, v_0, \kappa, \theta, \xi, \rho^{\text{lev}}$	Eigen-Cond Heston	Heston COS
∞	Full sample paths	Multi-asset Heston MC	Simulation

6.3 Limitations

1. **Assumes constant cross-asset correlation.** The C matrix is static. Stochastic correlation (dynamic $\rho_{ij}(t)$) would require extending the eigendecomposition to a time-varying framework.
2. **Heston CF may be slow for extreme parameters.** For $\xi > 2$ or $\kappa \approx 0$, the Heston CF oscillates and COS convergence slows.
3. **Calibration.** Each asset needs 5 Heston parameters ($\kappa, \theta, \xi, \rho^{\text{lev}}, v_0$). Multi-asset calibration from market smiles is a separate problem.

6.4 Open Questions

1. **Stochastic correlation.** Can the eigendecomposition be extended to a Wishart-type stochastic correlation model? This would require conditioning on both factor realizations AND correlation dynamics.
2. **Jump-diffusion extension.** The conditional CF under Merton/Kou jump-diffusion also factorizes. The same framework applies with the jump CF per scenario.
3. **Machine learning integration.** Use a neural network to learn the per-scenario Heston price as a function of $(F_i^{(q)}, v_i(0), \kappa_i, \theta_i, \xi_i, \rho_i^{\text{lev}}, K)$, bypassing the COS inversion entirely. This would combine the eigenvalue conditioning structure (fast, interpretable) with ML speed (microsecond per evaluation). The eigenvalue-conditional decomposition provides a natural feature space for the neural network, as discussed by Goudenège et al. (2020) [TODO:cite] in a related context.

7. Conclusion

The eigenvalue-conditional mixing framework extends naturally from GBM to stochastic volatility. The Mixture Collapse — the mathematical core of the approach — is distribution-agnostic (Lean-verified in `MixtureCollapse.lean`), requiring only linearity of finite sums and the law of iterated expectations. The extension replaces per-scenario Black-Scholes with per-scenario Heston COS pricing, preserving the eigendecomposition infrastructure and quadrature mixing unchanged.

The key mathematical contributions are: (i) the approximate factorization of the conditional basket characteristic function as a product of single-asset Heston CFs, with residual error bounded by the omitted eigenvalues (Section 1.3); (ii) the explicit conditional forward price formula that separates the cross-asset factor contribution from the per-asset stochastic volatility dynamics (Section 2.2); and (iii) the three-component error decomposition distinguishing correlation conditioning error, Heston CF truncation error, and quadrature error (Section 3.2).

To our knowledge, this is the first deterministic basket option pricing method under multi-asset stochastic volatility that exploits the spectral structure of cross-asset correlation, reducing the problem to a sequence of single-asset Heston pricings. Unlike Wishart-based approaches (Da Fonseca et al., 2008), which model the full covariance matrix as a stochastic process, our method decomposes the static cross-asset correlation and handles per-asset stochastic volatility through the single-asset Heston CF — a well-understood numerical building block.

Several directions remain open: a formal proof of the conditional variance reduction conjecture (Conjecture 1), extension to stochastic correlation models where C itself evolves in time, and integration with jump-diffusion dynamics where the conditional CF also factorizes. The computational framework is modular: any single-asset pricing model with an analytically tractable characteristic function can replace Heston without modifying the outer eigenvalue-conditioning and mixing infrastructure.

Lean Verification

The following results supporting this paper have been formally verified in Lean 4 (zero sorry, zero unverified axioms):

- **MixtureCollapse.lean:** The sum-interchange identity underlying the Mixture Collapse (Theorem 1). This is the core result: it verifies that the unconditional COS coefficients equal the quadrature-weighted sum of conditional coefficients, by commuting finite sums. This is the genuine formal verification that the mixing framework is distribution-agnostic.
- **ConditionalVariance.lean:** The variance decomposition identity and the bound $\tau_K \in [0, 1]$ for the explained variance ratio. These are elementary but foundational for the error analysis.
- **StochasticVolExtension.lean:** Contains `eigendecomposition_model_free` (elementwise equality implies matrix equality) and `stochastic_vol_mixing` (reflexivity). These verify structural properties of the formalism but are mathematically trivial — they confirm type-level consistency rather than deep mathematical content.

Scope of verification. The Lean proofs verify the algebraic infrastructure (sum interchange, variance decomposition) but do NOT formally verify the approximate CF factorization (Section 1.3), the conditional forward price derivation (Section 2.2), or Conjecture 1 (Section 3.1). These results rely on analytic arguments that are beyond the current scope of our formalization.

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

- Heston, S.L (1993). A closed-form solution for options with stochastic volatility. *Heston, S.L.*, 6(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
- Carr, Peter and Madan, Dilip (1999). Option Valuation Using the Fast Fourier. *Journal of Computational Finance*, 2(4), 61-73. DOI: 10.21314/jcf.1999.043
- Nagy, T. (2026). Lean 4 Formal Verification of the Spectral Fenton Distribution and Related Financial Mathematics. *Working paper*.
- Nagy, T. (2026). Pricing Basket Options via Eigenvalue-Conditional Black-Scholes Mixing. *Zenodo*. DOI: 10.5281/zenodo.18910542
- Da Fonseca, J., Grasselli, M., and Tebaldi, C (2008). A multifactor volatility Heston model. *Da Fonseca, J., Grasselli, M., and Tebaldi, C.*, 8(6). DOI: 10.1080/14697680701668418
- Merton, R.C (1976). Option pricing when underlying stock returns are discontinuous. *Merton, R.C.* DOI: 10.1016/0304-405x(76)90022-2
- Kou, S.G (2002). A jump-diffusion model for option pricing. *Kou, S.G.*, 48(8). DOI: 10.2139/ssrn.242367