Thesis Booklet: Essays on Derivative Pricing and Risk Management

Doctoral School of Economics and Business Informatics

Corvinus University of Budapest

Dániel Léber

Supervisors
Péter Csóka, PhD
Mihály Ormos, PhD

Budapest, 2026

1 Research Background

For both risk managers and portfolio managers, the continuous emergence of increasingly complex financial instruments poses a persistent challenge. While financial innovation expands the set of tools available for hedging, risk transfer, and portfolio construction, it also makes valuation and risk measurement more demanding. This is particularly true for exotic derivatives whose payoffs depend on the path of one or more underlying assets. Lookback, Asian, barrier, and rainbow options are well-known examples of such contracts, and their valuation typically requires methods that go beyond the standard closed-form results available for plain-vanilla claims. From a practical perspective, the problem is not merely to obtain derivative prices, but to obtain prices and risk measures with sufficient precision and speed for trading, hedging, portfolio rebalancing, and risk monitoring. For complex path-dependent and multi-asset products, the relevant dynamics are often too rich for simple analytical treatment, while standard numerical procedures may become costly when high accuracy is required. This creates a natural demand for methodologies that can deliver accurate valuation and practically useful risk measures in a computationally efficient way (Hull and Basu, 2016, Hull, 2023, Wilmott, 2013).

The challenge is compounded by the empirical properties of financial markets themselves. Asset returns are known to exhibit stylized facts such as volatility clustering, heavy tails, leverage effects, and intermittency, all of which require complex modelling and forecasting. At the same time, market dynamics are influenced not only by past prices but also by information flows and investor sentiment reflected in news, corporate language, internet message boards, and social media. This is particularly important because financial institutions must forecast risk measures with a high degree of accuracy, while also meeting increasingly stringent regulatory requirements.

In this context, volatility forecasts play a central role in Value-at-Risk forecasting used for internal risk control, capital allocation, and regulatory backtesting (Cont, 2001).

This dissertation comprises three main chapters. The first two develop numerical and semi-analytical Hamiltonian methods for pricing single- and multi-asset derivatives, while the third examines sentiment-augmented models for forecasting volatility and Value-at-Risk.

2 Research Questions and Hypotheses

These challenges motivate the present dissertation, which examines how analytical, numerical, and data-driven methods can contribute to high-precision valuation and risk measurement of path-dependent derivatives, as well as to sentiment-based volatility forecasting. More specifically, the dissertation seeks answers to the following questions:

1. How can pseudo-Hermitian Hamiltonian simulations be turned into practical high-precision algorithms for pricing and risk measurement of path-dependent options?

2. How can this methodology be generalised to the case of multiple correlated underlying assets in order to price multi-asset path-dependent derivatives?

3. How can volatility and Value-at-Risk forecasting be improved by combining GARCH-type models, recurrent neural networks, and sentiment information extracted from news and social media?

The first two research questions give rise to one methodological hypothesis each, while the third research question is examined through two empirical hypotheses concerning the predictive content of sentiment information and its contribution to forecasting performance. Accordingly, the dissertation formulates the following hypotheses:

H1. It is hypothesised that, within the examined single-asset Black–Scholes–Merton settings, the pseudo-Hermitian Hamiltonian spectral method will attain numerical accuracy at least comparable to that of the Crank–Nicolson, binomial-tree, and Monte Carlo benchmark methods for option valuation, with particularly strong performance for quantities that depend sensitively on the reconstructed Arrow–Debreu state-price density. On sufficiently fine grids, the shift-invert Lánczos approximation is expected to retain essentially the accuracy of the full-spectral solution while exhibiting a more favourable empirical runtime profile.

H2. It is hypothesised that the pseudo-Hermitian similarity transformation considered for the single-asset problem in H1 can be generalised into a Hermitian framework for pricing simple and path-dependent options on multiple correlated underlying assets, thereby enabling the derivation of semi-analytical pricing solutions and efficient approximations for path-dependent multi-asset derivatives.

H3. It is hypothesised that positive and negative sentiment indices derived from news and Twitter/X contain statistically significant predictive information for the conditional variance of daily returns among the S&P 500 stocks included in the empirical sample.

H4. It is hypothesised that augmenting the GARCH and GARCH–LSTM-based models with sentiment information will reduce out-of-sample realised-volatility forecast errors for stocks included in the empirical sample relative to the corresponding models without sentiment. The resulting improvements in volatility forecasts are, in turn, expected to improve the performance of Value-at-Risk forecasts, as assessed using standard Value-at-Risk backtests.

2.1 Design of the chapters

The dissertation is organised into three substantive chapters. Table 1 presents their exact titles, the corresponding research questions, and the applied methodologies.

Table 1: Research design of the dissertation

Chapter Title	**Research Question**	**Applied Methodology**
Efficient High Precision Pseudo-Hermitian Hamiltonian Simulation of Path-Dependent Options	How can pseudo-Hermitian Hamiltonian simulations be turned into practical high-precision algorithms for pricing and risk measurement of path-dependent options?	Pseudo-Hermitian similarity transformation, Hamiltonian spectral simulation, shift-invert Lánczos method, Arrow–Debreu state-price-density reconstruction, and benchmarking against Crank–Nicolson, binomial-tree, and Monte Carlo methods.
Quantum Mechanical Approach to Pricing Multi-Asset Path-Dependent Options	How can the pseudo-Hermitian methodology be generalised to the case of multiple correlated underlying assets in order to price multi-asset path-dependent derivatives?	Multi-asset pseudo-Hermitian similarity transformation, momentum-space wave functions and superposition, and a multi-particle-in-a-box construction for the two-asset double knock-out worst-of rainbow option.
The Sentiment Augmented GARCH-LSTM Hybrid Model for Value-at-Risk Forecasting	How can volatility and Value-at-Risk forecasting be improved by combining GARCH-type models, recurrent neural networks, and sentiment information extracted from news and social media?	Linear and neural-network-based nonlinear Granger causality tests, GARCH(1,1) and GARCH–LSTM forecasting, volatility-error measures, and Value-at-Risk backtesting.

3 Scientific Results and Thesis Statements

3.1 Chapter 2. Efficient High Precision Pseudo-Hermitian Hamiltonian Simulation of Path-Dependent Options

T1. Within the examined single-asset Black–Scholes–Merton settings, the pseudo-Hermitian Hamiltonian spectral method reconstructs the Arrow–Debreu state-price density with high numerical accuracy and is fully competitive with the Crank–Nicolson, binomial-tree, and Monte Carlo benchmark methods for option valuation, with particularly strong performance for density-sensitive quantities. On sufficiently fine grids, the shift-invert Lánczos approximation reproduces the full-spectral estimates with essentially unchanged accuracy while exhibiting a more favourable empirical runtime.

Let $x = \ln S$ and $\mu = r - \sigma^2/2$. In log-price coordinates, the one-asset pricing Hamiltonian is

$$\widehat{H}_{\mathrm{BS}} = r - \mu \frac{d}{dx} - \frac{\sigma^2}{2}\frac{d^2}{dx^2}. \tag{1}$$

The state-dependent transformation

$$\widehat{\rho}\widehat{H}_{\mathrm{BS}}\widehat{\rho}^{-1} = \widehat{h}_{\mathrm{BS}}, \qquad \widehat{\rho} = \exp\left[-\left(\frac{1}{2} - \frac{r}{\sigma^2}\right)x\right] \tag{2}$$

removes the first derivative and gives

$$\widehat{h}_{\mathrm{BS}} = -\frac{\sigma^2}{2}\frac{d^2}{dx^2} + \frac{1}{2\sigma^2}\left(\frac{\sigma^2}{2} + r\right)^2, \qquad \widehat{\eta} = \widehat{\rho}^2. \tag{3}$$

Here $\rho(x) = \exp[-(1/2 - r/\sigma^2)x]$ denotes the scalar multiplier associated with the operator $\widehat{\rho}$. Economically, $\widehat{\rho}$ is an invertible state-dependent reweighting that cancels the first-order risk-neutral log-return drift in the transformed representation. The drift's directional effect is retained in the endpoint factors $\rho(x_t)^{-1}\rho(x_T)$, while the constant term of $\widehat{h}_{\mathrm{BS}}$ carries discounting together with the drift correction, so the original Black–Scholes–Merton valuation is unchanged.

For time to maturity $\tau = T - t$, the Arrow–Debreu state-price density is the coordinate-space kernel of the pricing semigroup,

$$\pi_{t,T}(x_t, x_T) = \langle x_t | e^{-\tau \widehat{H}_{\mathrm{BS}}} | x_T \rangle = \rho(x_t)^{-1}\rho(x_T)\langle x_t | e^{-\tau \widehat{h}_{\mathrm{BS}}} | x_T \rangle, \quad (4)$$

where $x_t = \ln S_t$ and $x_T = \ln S_T$. A terminal payoff g_T is valued over the relevant log-price domain D by

$$g(x_t, \tau) = \int_D \pi_{t,T}(x_t, x_T) g_T(x_T)\, dx_T. \quad (5)$$

A Green function and the Arrow–Debreu kernel play analogous roles. Each propagates a point source through the governing linear operator. In the spectral representation, each eigenmode is weighted by $e^{-\epsilon_n \tau}$. Since higher modes are exponentially suppressed as time to maturity increases, their contribution becomes negligible beyond a moderate index (Davydov and Linetsky, 2003, Linetsky, 2007). The shift-invert Lánczos Krylov strategy therefore targets only the eigenpairs required to satisfy a prescribed spectral-tail tolerance, avoiding full diagonalisation while retaining the stability benefits of an orthogonal eigenbasis.

To evaluate the numerical accuracy of the pricing algorithms, Table 2 compares the plain-vanilla option values obtained from the different numerical methods with the analytical Black–Scholes–Merton price. Among all methods considered, the Hamiltonian spectral approach achieves the smallest absolute error, followed very closely by the Hamiltonian Lánczos spectral method. The additional discrepancy introduced by the Lánczos approximation is negligible and remains essentially at the level of machine precision relative to the simple Hamiltonian spectral solution. The next most accurate result is provided by the Crank–Nicolson method, followed by the binomial tree method, while the Monte Carlo simulation exhibits the largest deviation from the analytical benchmark.

Table 2: Plain-vanilla call valuation against the analytical BSM price

Method	Price	Absolute error	Relative error
Analytical BSM	1.9520993	–	–
Monte Carlo	1.9647231	1.2624×10^{-2}	6.4668×10^{-3}
Binomial tree	1.9491788	2.9205×10^{-3}	1.4961×10^{-3}
Crank–Nicolson	1.9526431	5.4378×10^{-4}	2.7856×10^{-4}
Full spectral	1.9524891	3.8982×10^{-4}	1.9969×10^{-4}
Shift-invert Lánczos	1.9524891	3.8982×10^{-4}	1.9969×10^{-4}

For the double knock-out barrier case, the Hamiltonian spectral and Crank–Nicolson methods generate highly consistent state-price densities and state-price-implied risk measures, while the Lánczos approximation remains nearly indistinguishable from the full spectral solution. Relative to the full spectral solution, the Lánczos density has an L^1 discrepancy of 2.148×10^{-12}. The Crank–Nicolson discrepancy is 1.668×10^{-7}. The Hamiltonian spectral, Lánczos, and Crank–Nicolson constructions imply the same knock-out probability to four decimal places, which is 10.0979%, and the same survival probability of 89.9021%. The Lánczos estimates of the double knock-out Delta and surviving digital value differ from the full spectral reference by approximately 2.21×10^{-12} and 7.39×10^{-13} in relative terms.

For the plain-vanilla European call, the Black–Scholes–Merton closed form is used to assess approximation error and runtime across the competing numerical approaches. A more informative perspective is provided by the runtime-accuracy diagram shown in Figure 1, which plots the absolute pricing error against runtime on logarithmic axes. The conventional dense spectral implementation already yields

Figure 1: Runtime versus absolute error against the analytical Black–Scholes–Merton price

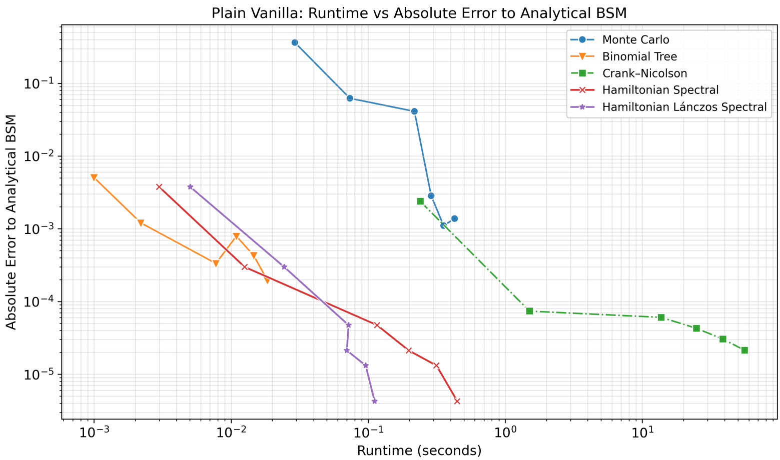

low errors at comparatively short runtimes, whereas the Hamiltonian Lánczos spectral method becomes increasingly advantageous as the resolution increases. Beyond the first low-resolution cases, it becomes increasingly favourable and eventually attains essentially the same accuracy as the full spectral method, while requiring less computational time.

H1 is supported within the examined numerical setting. The results of Chapter 2 show that the pseudo-Hermitian Hamiltonian framework is fully competitive with the Crank–Nicolson, binomial-tree, and Monte Carlo benchmarks, with particularly strong performance for quantities that depend sensitively on the Arrow–Debreu state-price density. On sufficiently fine grids, the shift-invert Lánczos method retains essentially the accuracy of the full-spectral solution while exhibiting more favourable empirical runtime scaling.

3.2 Chapter 3. Quantum Mechanical Approach to Pricing Multi-Asset Path-Dependent Options

T2. The pseudo-Hermitian similarity transformation applied in the single-asset framework of T1 is generalisable to correlated multi-asset Black–Scholes–Merton pricing problems and maps the corresponding non-Hermitian pricing Hamiltonian to a Hermitian representation. For European terminal payoffs, the resulting free-particle formulation recovers the conventional multi-asset Black–Scholes–Merton pricing kernel and thereby enables the valuation of standard multi-asset European options. For the examined two-asset double knock-out worst-of rainbow option with asset-specific barriers, the particle-in-a-box construction yields a computationally tractable semi-analytical approximation.

Chapter 3 generalises the single-asset framework of T1 to multiple correlated underlying assets. Its main contribution is the derivation of the pseudo-Hermitian Hamilton operator of the multi-asset Black–Scholes–Merton model. Let $\mathbf{x}$ contain the logarithms of the N underlying prices, let $\mathbf{\Sigma}$ be their positive-definite instantaneous covariance matrix, and let $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_N)^{\mathsf{T}}$ with $\mu_i = r - \sigma_i^2/2$. Let $\widehat{\nabla} = (\partial/\partial x_1, \ldots, \partial/\partial x_N)^{\mathsf{T}}$. The correlated multi-asset BSM Hamiltonian can be written as

$$\widehat{H}_{\mathrm{BSM}} = r - \boldsymbol{\mu}^{\mathsf{T}}\widehat{\nabla} - \frac{1}{2}\widehat{\nabla}^{\mathsf{T}}\mathbf{\Sigma}\widehat{\nabla}. \tag{6}$$

As in the single-asset formulation of Chapter 2, the pseudo-Hermitian similarity transformation is

$$\begin{aligned} \widehat{h} &= \widehat{\gamma}\widehat{H}_{\mathrm{BSM}}\widehat{\gamma}^{-1}, \\ \widehat{\gamma} &= \exp\!\big(\mathbf{x}^{\mathsf{T}}\mathbf{\Sigma}^{-1}\boldsymbol{\mu}\big), \quad \widehat{\eta} = \widehat{\gamma}^2 \end{aligned} \tag{7}$$

It cancels the vector of risk-neutral log-return drifts from the first-order terms and gives the Hermitian representation

$$\widehat{h} = -\frac{1}{2}\widehat{\nabla}^{\mathsf{T}}\boldsymbol{\Sigma}\widehat{\nabla} + \frac{1}{2}\boldsymbol{\mu}^{\mathsf{T}}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} + r. \tag{8}$$

The drifts are retained in the endpoint $\widehat{\gamma}$ factors, while the zero-order term carries discounting together with the drift correction, so the original BSM valuation is unchanged. In the single-asset case, both the $\widehat{\gamma}$ operator and the $\widehat{h}$ Hermitian BSM Hamilton operator coincide with the operators defined by Jana and Roy (2012). In the multi-asset case, there are also partial cross-derivatives, which, according to Young's theorem, result in a symmetric operator. If the coefficients of partial cross-derivatives are real in Hamilton operators, then these operators are also Hermitian.

Using Fourier transformation, we seek the solution for the general multi-asset option in the form of a momentum-space wave function, analogous to determining the wave function of a free particle in quantum mechanics. Momentum-space wave functions and superposition are used to price general European multi-asset options, for which the conventional multi-asset Black–Scholes–Merton pricing kernel is recovered.

A significant class of financial derivatives consists of path-dependent, limit-price options that include instruments on the basis of multiple underlying assets. These options often incorporate specific price thresholds. We present the pricing of a two-asset double knock-out rainbow option that incorporates two individual limit prices for the underlying assets. This option value becomes zero as soon as one of the two underlying assets breaches its respective barrier level.

The Hermitian BSM Hamiltonian operator also contains partial cross-derivatives, which can be eliminated by diagonalisation to achieve greater ease of solution. We introduce y_1 and y_2 according to

$$\begin{aligned} y_1 &= x_1 \cos\theta + x_2 \sin\theta \\ y_2 &= -x_1 \sin\theta + x_2 \cos\theta \end{aligned} \tag{9}$$

where θ ensures that the cross-derivative operator vanishes from the corresponding Hamilton operator. As a result of diagonalisation, the Hermitian BSM Hamiltonian operator can be expressed as

$$\widehat{h} = -\frac{1}{2}\left(\lambda_1 \frac{\partial^2}{\partial y_1^2} + \lambda_2 \frac{\partial^2}{\partial y_2^2}\right) + \widehat{V}. \tag{10}$$

Here λ_1 and λ_2 are the eigenvalues of $\mathbf{\Sigma}$. In quantum mechanics, the boundary conditions of a particle are inherently encoded in the potential energy operator $\widehat{V}$, as it is directly influenced by the position operator. One specific realisation of this principle is the infinite potential well, a system in which the particle is confined and cannot escape the well. Consequently, the path dependence of various limit-price options can be effectively addressed through their potential energy operator $\widehat{V}$.

The numerical illustrations use $T = 3$ months, $r = 5\%$, $\sigma_1 = 40\%$, and $\sigma_2 = 20\%$. With lower and upper limit prices of 90 and 120, a strike price of 100, and a correlation of 0.05, the worst-of rainbow option value reaches about 0.24 near $(S_1, S_2) = (104, 107)$. When the correlation is 0.6, the maximum option value is about 0.66 near $(102, 109)$ and the high-value region becomes more elongated. It is also possible to apply different limit prices to the underlying assets, allowing for increased flexibility of the conditions attached to the underlying assets. Figure 2 uses a strike price of 110 and a correlation of 0.6.

Chapter 3 demonstrates that the pseudo-Hermitian similarity transformation can be generalised to correlated multi-asset Black–Scholes–Merton pricing problems and that it maps the corresponding

Figure 2: The heatmap of the two-asset double knock-out worst-of rainbow option prices as a function of the prices of the two underlying assets. S_1 has limit prices between 100 and 140, while S_2 has limit prices between 90 and 130. The correlation between the two underlying assets is assumed to be 0.6.

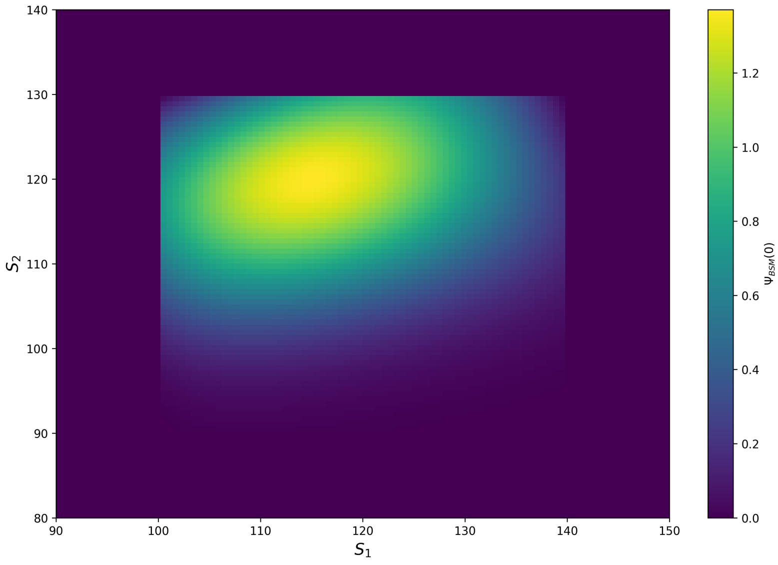

non-Hermitian pricing Hamiltonian to a Hermitian representation. For European terminal payoffs, the resulting free-particle formulation recovers the conventional multi-asset Black–Scholes–Merton pricing kernel, thereby enabling the valuation of standard multi-asset European options. For the examined two-asset double knock-out worst-of rainbow option with asset-specific barriers, the particle-in-a-box construction approximates the transformed barrier problem and yields a computationally tractable semi-analytical approximation method. Taken together, these results support H2.

The results underlying T2 were published in a single-author article in *Physica A: Statistical Mechanics and its Applications* (Léber, 2025).

3.3 Chapter 4. The Sentiment Augmented GARCH-LSTM Hybrid Model for Value-at-Risk Forecasting

T3. Lagged positive and negative sentiment indices derived from news and Twitter/X contain statistically significant predictive information for the conditional variance of daily returns among substantial proportions of the examined S&P 500 stocks. The considerably higher detection rates of the nonlinear tests show that nonlinear specifications capture a substantial part of the predictive information contained in sentiment that remains undetected by the standard linear test.

The data used for the analysis were downloaded exclusively from the Bloomberg Terminal. Individual stocks in the S&P 500 stock index for March 2024 are used for the analysis. Individual stocks' adjusted close prices, Twitter/X publication counts, Twitter/X positive and negative sentiment counts, news publication counts, and news positive and negative sentiment counts are downloaded at a daily frequency. The impact of the negative and positive sentiment indices is examined separately because, although these indicators have a negative correlation, they cannot be perfectly expressed because of the inclusion of a nonnegligible amount of neutral sentiments.

The linear Granger causality test is conducted only in cases where both the variance and the given standardised sentiment index are stationary at the 5% significance level. The statistical tests are performed between the variance and the different standardised sentiment indices with lags ranging from 1 to 10. To determine whether there is a nonlinear relationship between the variance and different sentiment indices, we evaluate the nonlinear Granger causality approach proposed by Rosol et al. (2022).

H3 is partially supported because statistically significant predictive ability was not found for all examined stocks. At the first lag and the 5% significance level, the nonlinear tests detected significant predictive relationships for 83%, 81%, 39%, and 37% of the examined stocks for negative Twitter/X, positive Twitter/X, negative news, and positive news sentiment, respectively. The corresponding linear detection rates were 27%, 9%, 6%, and 6%. Thus, the results indicate that nonlinear specifications capture a substantial part of the predictive information contained in sentiment that remains undetected by linear tests. The corresponding first-lag detection rates are summarised in Table 3.

Table 3: First-lag predictive detection rates at the 5% significance level

Sentiment measure	Linear	Nonlinear	Difference
Negative Twitter/X	27%	83%	56 pp
Positive Twitter/X	9%	81%	72 pp
Negative news	6%	39%	33 pp
Positive news	6%	37%	31 pp

3.3.1 Performance Evaluation and Backtesting of Predictive Models

T4. Augmenting the GARCH and GARCH–LSTM-based models with sentiment information reduces out-of-sample realised-volatility forecast errors for a majority of the examined stocks under each of MSE, RMSPE, and MAE. These forecasting gains are accompanied by stronger overall Value-at-Risk backtesting results across the coverage, duration, and joint tests, with the sentiment-augmented specifications attaining higher non-rejection rates in five of the six model-by-test comparisons.

The chapter uses the one-day ahead expanding rolling window technique to forecast one-day Value-at-Risk values at the 2.5% significance level using four models. The first model is the traditional GARCH(1,1) model. The second model is the sentiment-augmented GARCH(1,1) model. The third and fourth models are LSTM neural network models, both of which are based on GARCH(1,1) forecasts. The first neural network consists solely of contemporaneous and lagged volatility forecasts from the GARCH(1,1) model. The second neural network is augmented with different lagged sentiment index values. The two neural networks differ solely in their input variables.

Volatility forecasts are evaluated by mean squared error (MSE), root mean square percentage error (RMSPE), and mean absolute error (MAE). Value-at-Risk forecasts are assessed through Kupiec's unconditional coverage test, the Christoffersen–Pelletier duration test, and their joint test at the 5% significance level (Campbell, 2005, Christoffersen and Pelletier, 2004). Table 4 shows the proportion of stocks for which the augmented model has the smaller forecast error and the percentage-point changes in the proportion of stocks for which the corresponding Value-at-Risk null hypothesis is not rejected.

H4 is also partially supported because sentiment augmentation improved forecast performance for a majority, but not all of the

Table 4: Pairwise effects of sentiment augmentation on volatility errors and VaR backtests

Panel A. Stocks with lower forecast error			
Comparison	MSE	RMSPE	MAE
Sentiment-augmented GARCH(1,1) vs GARCH(1,1)	61.91%	75.96%	73.19%
Sentiment-augmented GARCH(1,1)–LSTM vs GARCH(1,1)–LSTM	69.57%	79.36%	76.38%
Panel B. Change in Value-at-Risk non-rejection rate			
Comparison	Coverage	Duration	Joint
Sentiment-augmented GARCH(1,1) vs GARCH(1,1)	+14.68 pp	+6.60 pp	+5.32 pp
Sentiment-augmented GARCH(1,1)–LSTM vs GARCH(1,1)–LSTM	−6.38 pp	+20.00 pp	+11.06 pp

examined stocks. Relative to GARCH, the sentiment-augmented GARCH model produced lower errors for 61.91%, 75.96%, and 73.19% of the stocks under MSE, RMSPE, and MAE, respectively. Relative to the corresponding GARCH–LSTM-based benchmark, the sentiment-augmented model improved these measures for 69.57%, 79.36%, and 76.38% of the stocks. The Value-at-Risk results also improved overall, as sentiment augmentation generally increased the non-rejection rates across the applied backtests. For the GARCH specifications, improvements occurred under the coverage, duration, and joint tests, whereas for the GARCH–LSTM-based specifications the gains were concentrated in the duration and joint tests. The improvement was therefore substantial overall, although it was not uniform across all model specifications and backtests.

The results underlying T3 and T4 were obtained jointly with Balázs Egyed and published in *Computational Economics* (Léber and Egyed, 2025).

4 Main References

Campbell, S.D., 2005. A Review of Backtesting and Backtesting Procedures. Finance and Economics Discussion Series 2005-21. Board of Governors of the Federal Reserve System (U.S.). doi:https://doi.org/10.17016/FEDS.2005.21.

Christoffersen, P., Pelletier, D., 2004. Backtesting value-at-risk: A duration-based approach. Journal of Financial Econometrics 2, 84–108. doi:https://doi.org/10.1093/jjfinec/nbh004.

Cont, R., 2001. Empirical properties of asset returns: stylized facts and statistical issues. Quantitative finance 1, 223. doi:https://doi.org/10.1080/713665670.

Davydov, D., Linetsky, V., 2003. Pricing Options on Scalar Diffusions: An Eigenfunction Expansion Approach. Operations Research 51, 185–209. doi:https://doi.org/10.1287/opre.51.2.185.12782.

Hull, J.C., 2023. Risk management and financial institutions. John Wiley & Sons.

Hull, J.C., Basu, S., 2016. Options, Futures, and Other Derivatives. Pearson Education India.

Jana, T., Roy, P., 2012. Pseudo Hermitian formulation of the quantum Black–Scholes Hamiltonian. Physica A: Statistical Mechanics and its Applications 391, 2636–2640. doi:https://doi.org/10.1016/j.physa.2011.12.012.

Léber, D., 2025. Quantum mechanical approach to pricing multi-asset path-dependent options. Physica A: Statistical Mechanics and its Applications 675, 130644. doi:https://doi.org/10.1016/j.physa.2025.130644.

Léber, D., Egyed, B., 2025. The sentiment augmented GARCH-LSTM hybrid model for value-at-risk forecasting. Computational Economics , 1–41 doi:https://doi.org/10.1007/s10614-025-11042-8.

Linetsky, V., 2007. Spectral Methods in Derivatives Pricing, in: Birge, J.R., Linetsky, V. (Eds.), Financial Engineering. Elsevier. volume 15 of *Handbooks in Operations Research and Management Science*, pp. 223–299. doi:https://doi.org/10.1016/S0927-0507(07)15006-4.

Rosol, M., Młyńczak, M., Cybulski, G., 2022. Granger causality test with nonlinear neural-network-based methods: Python package and simulation study. Computer Methods and Programs in Biomedicine 216, 106669. doi:https://doi.org/10.1016/j.cmpb.2022.106669.

Wilmott, P., 2013. Paul Wilmott Introduces Quantitative Finance. The Wiley Finance Series, Wiley.

Publications Underlying the Dissertation

Léber, D., 2025. Quantum mechanical approach to pricing multi-asset path-dependent options. Physica A: Statistical Mechanics and its Applications 675, 130644. doi:https://doi.org/10.1016/j.physa.2025.130644.

Léber, D., Egyed, B., 2025. The sentiment augmented GARCH-LSTM hybrid model for value-at-risk forecasting. Computational Economics , 1–41 doi:https://doi.org/10.1007/s10614-025-11042-8.