

RISK QUANTIFICATION AND BAYESIAN CALIBRATION: ASSESSING TAIL RISK ACROSS MARKET REGIMES

Riskin Nicel Ölçümü ve Bayesçi Kalibrasyon: Piyasa Rejimleri Arasında Kuyruk
Riskinin Değerlendirilmesi

Hao Jin^{*}, Hammad Siddiqi^{**}, Sajid Anwar^{***} & Jingdong Huang^{****}

Abstract

Post-pandemic volatility clustering highlights the gap between static governance and quantitative risk measurement. This paper proposes a synthesized framework utilizing complementary components of classical risk measurement and Bayesian inference. First, applying Lagrangian optimization, we illustrate the sensitivity of static mean-variance allocations to regime-specific sample moments. Second, we introduce Expected Shortfall as a complementary theoretical measure and GARCH specifications to account for time-varying conditional variance. Crucially, we complement standard VaR backtesting with a Beta-Binomial conjugate prior framework to calibrate risk exposures. Utilizing S&P 500 daily returns, we demonstrate how this mechanism quantifies the uncertainty surrounding the VaR violation probability. Relative to the 1.0% expectation, the traditional 21-Day Rolling VaR yields a Bayesian breach-rate adjustment factor of 2.40x, whereas the GARCH(1,1) specification lowers this to 1.80x. Ultimately, by bridging the gap between ex-ante predictions and realized shocks, this Bayesian recalibration loop may provide a data-driven approach to ex-post risk calibration.

Keywords:

Portfolio Optimization,
Expected Shortfall
(ES), Value-at-Risk
(VaR); Bayesian
Calibration; Tail Risk

JEL Codes:

G32, C11, G11.

Öz

Pandemi sonrası dönemde volatilité kümelenmesinin belirginleşmesi, statik yönetim yaklaşımları ile nicel risk ölçümü arasındaki boşluğu ortaya çıkarmıştır. Bu çalışma, klasik risk ölçüm yöntemleri ile Bayesçi çıkarımın birbirini tamamlayan bileşenlerini bir araya getiren bütünlük bir çerçeve önermektedir. İlk olarak, Lagrange optimizasyonu kullanılarak statik ortalama-varyans portföy tahsislerinin rejime özgü örneklem momentlerine duyarlılığı ortaya konulmaktadır. İkinci olarak, tamamlayıcı bir teorik risk ölçütü olarak Beklenen Kayıp (Expected Shortfall) ele alınmakta ve zamanla değişen koşullu varyansı modellemek amacıyla GARCH spesifikasyonlarından yararlanılmaktadır. Çalışmanın temel katkısı, standart VaR geri testlerinin, risk maruziyetlerinin kalibrasyonu amacıyla Beta-Binomial eşlenik önsel çerçevesiyle tamamlanmasıdır. S&P 500 günlük getirileri kullanılarak, bu mekanizmanın VaR ihlal olasılığına ilişkin belirsizliği nicel olarak ortaya koyduğu gösterilmektedir. %1,0'luk teorik ihlal oranı ile karşılaştırıldığında, geleneksel 21 günlük hareketli VaR yaklaşımı Bayesçi ihlal oranı düzeltme katsayısını 2,40 kat olarak üretirken, GARCH(1,1) spesifikasyonu bu katsayıyı 1,80 kata düşürmektedir. Sonuç olarak, önsel tahminler ile gerçekleşen şoklar arasındaki boşluğu gidermeyi amaçlayan bu Bayesçi yeniden kalibrasyon döngüsü, gerçekleştirmeler sonrasında risk ölçümlerinin veri odaklı biçimde yeniden kalibre edilmesine yönelik bir yaklaşım sunmaktadır.

Anahtar Kelimeler:

Portföy
Optimizasyonu,
Beklenen Kayıp, Riske
Maruz Değer (VaR),
Bayesçi Kalibrasyon,
Kuyruk Riski

JEL Kodları:

G32, C11, G11.

* Sessional Lecturer, University of the Sunshine Coast, School of Business and Creative Industries, Australia, hjin1@usc.edu.au

** Senior Lecturer, Dr., University of the Sunshine Coast, School of Business and Creative Industries, Australia, hsiddiqui@usc.edu.au

*** Prof. Dr., University of the Sunshine Coast, School of Business and Creative Industries, Australia, sanwar@usc.edu.au

**** Postgraduate Researcher, Tibet University, College of Finance and Economics, China, hjd5432@163.com (Corresponding Author)

Received Date (Makale Geliş Tarihi): 01.03.2026 Accepted Date (Makale Kabul Tarihi): 24.09.2026

This article is licensed under Creative Commons Attribution 4.0 International License.



1. Introduction

The stability of the modern financial system depends in part on the robustness of equity market portfolios and institutional risk management frameworks. Geopolitical shocks and the liquidity disruptions associated with the COVID-19 pandemic illustrate that market risk can propagate across assets and institutions, creating substantial challenges for portfolio construction and risk measurement. Such episodes can affect the risk exposures of individual institutions and the broader investment environment.

In the current market environment, financial institutions face a dual challenge. On one hand, the volume of financial assets has expanded substantially, driven by monetary easing and financial deepening. On the other hand, the characteristics of financial risk have evolved. Financial returns often exhibit non-normal features, including fat tails and skewness. Traditional qualitative management, while essential for governance culture, lacks the precision to handle these non-linear threats. A purely quantitative approach may also be limited when disconnected from economic context. Mathematical models that assume perfect liquidity or constant correlations may perform poorly during periods of market stress. Therefore, effective risk management requires an approach that combines theoretical risk identification with quantitative risk measurement.

However, standard quantitative models may not fully represent the uncertainty surrounding the Value-at-Risk (VaR) violation probability. Such models frequently ignore the uncertainty, leading to a “model risk” where model-implied risk estimates may diverge from realized exceedance performance across market periods. This creates a critical gap between theoretical risk measurement and empirical reality.

To bridge the gap between theoretical risk measurement and empirical reality, this study develops a risk-measurement framework combining complementary components for tail-risk calibration across market regimes. Specifically, this paper makes three core contributions to the existing literature. First, we show that both the 21-Day Rolling Window and conditional GARCH(1,1) VaR models exhibit violation frequencies consistent with inadequately captured tail risk across the pre- and post-pandemic periods. Second, we implement a one-day VaR backtesting framework. Utilizing Kupiec and Christoffersen tests, we assess differences in the serial dependence of VaR violations across the pre- and post-pandemic periods. Third, based on these empirical diagnostics, we construct a Beta-Binomial Bayesian conjugate framework. This model explicitly quantifies the uncertainty surrounding the VaR violation probability and derives a Bayesian breach-rate adjustment factor for institutional risk monitoring.

The remainder of this paper is organized as follows. Section 2 reviews the literature. Section 3 outlines the quantitative methodology, encompassing the portfolio optimization, VaR, Expected Shortfall (ES) as a complementary theoretical measure, and the Bayesian updating framework. Section 4 presents a foundational empirical application on asset allocation shifts across market regimes. Section 5 delivers the empirical results, detailing the S&P 500 daily backtesting results, statistical diagnostic tests, and the dynamic Bayesian breach-rate calibration. Section 6 discusses the implications of the proposed risk-management framework, and Section 7 concludes.

2. Literature Review and Theoretical Taxonomy

2.1. The Evolution of Risk Management Theory

The conceptual architecture of modern financial risk management has undergone a profound transformation, moving from static variance-based frameworks to time-varying, tail-sensitive paradigms (Naifar, 2025). The intellectual foundation was established by Markowitz (Fabozzi et al., 2008), whose Mean-Variance Optimization introduced the principle of diversification as a formal mathematical objective. However, the subsequent decades revealed that traditional implementations of Modern Portfolio Theory (MPT) may not fully capture asymmetric and tail behavior in equity returns, which can generate substantial losses during extreme market events (Kan and Lassance, 2025). In response to the liquidity crises of the 1990s, VaR emerged as the industry's first standardized quantification of downside potential.

While VaR was a cornerstone of the Basel II regulatory framework, its lack of sub-additivity—a core requirement for coherent risk measures as defined by Artzner et al. (1999)—motivated the development and adoption of alternative tail-risk measures (Acerbi and Tasche, 2002). Recent critiques, such as those by Zhu et al. (2025), argue that VaR frequently underestimates risks during periods of high persistent volatility, failing to capture the “integrated tail shape” of modern assets. Research into extreme tail behavior by Le (2024) further suggests that risk persistence in financial series necessitates more robust estimation methods. This theoretical deficiency contributed to the transition toward ES in the Basel III Fundamental Review of the Trading Book (FRTB). Unlike VaR, ES provides an average measure of losses within the distribution's tail, satisfying the sub-additivity axiom and providing information on the severity of tail losses during periods of extreme market dislocation. Contemporary research by Kuang (2024) further supports this shift, demonstrating through Extreme Value Theory (EVT) that ES is significantly more reliable than VaR in assessing asset classes prone to sudden rallies in high-volatility regimes.

2.2. Volatility Clustering and VaR Backtesting Challenges

Financial markets are inherently characterized by volatility clustering and shifts across market regimes, particularly during global crises. To capture these non-linear dynamics, recent literature frequently employs the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) family of models as a baseline econometric tool. For instance, Arfaoui and Yousaf (2022) utilized a VAR asymmetric BEKK-GARCH model to examine volatility spillovers across commodity and equity markets during the COVID-19 pandemic. Similarly, Maharana et al. (2024) and Maharana et al. (2025) applied DCC- and BEKK-GARCH frameworks to analyze volatility transmission in Indian stock markets under global shocks. Other studies have incorporated these conditional models to assess causal relationships in mining stocks (Nasrudin et al., 2024), evaluate macroeconomic impacts on GDP (Jui et al., 2024), and investigate safe-haven assets (Dimitriadis et al., 2025). In the context of risk measurement, Huang et al. (2024) employed an ARMA-GARCH-VaR approach to estimate cryptocurrency tail risks during severe market fluctuations.

Despite the widespread application of VaR, validating these risk models during turbulent periods presents significant backtesting challenges. Kakade et al. (2022) investigated VaR forecasting during the 2008 and 2020 crises, emphasizing that even standard GARCH models

frequently fail to efficiently capture volatility during extreme market downturns. To rigorously validate risk models against such failures, recent empirical research relies heavily on formal diagnostic tools—predominantly Kupiec’s Proportion of Failures (POF) test for unconditional coverage and Christoffersen’s independence test to evaluate the serial independence of VaR exceptions. Employing these standardized metrics, Muhtarulloh et al. (2026) and Mufaridho et al. (2026) evaluated model accuracy across currency options and stock portfolios, respectively. While comparative studies by Lam (2025) suggest that conditional approaches can provide relative improvements over traditional historical methods, existing models broadly struggle with residual tail exceptions. Notably, Mallela and Gupta (2024) examined unweighted static VaR models for cryptocurrencies, revealing significant inadequacies in accounting for the clustering of VaR violations under extreme market conditions.

2.3. Bayesian Approaches to VaR Uncertainty

A critical weakness of standard VaR and ES implementations is that they may not explicitly represent uncertainty surrounding the VaR violation probability. Recent scholarship by Babic (2019) suggests that transitioning to ES alone is insufficient to protect institutional capital if uncertainty in risk estimation is left unaddressed. As noted by Balakrishnan and Ertan (2021), the “recognition lag” frequently observed during sudden market drawdowns is associated with delayed updating of risk estimates.

To mitigate this vulnerability, there is an emerging interest in Bayesian inferential frameworks that explicitly account for uncertainty in risk parameters and risk estimates (Beckmann et al., 2022; Feldkircher et al., 2022). Unlike traditional frequentist models, Bayesian logic treats risk parameters as probability distributions that are continuously updated in response to realized market shocks. Recent empirical studies report improved performance for Bayesian approaches in volatile environments; for instance, Carriero et al. (2024) show that Bayesian VaR and ES estimation significantly outperforms frequentist methods across different market regimes. While advanced Bayesian methods have been explored theoretically in risk management—such as the application of Markov Chain Monte Carlo (MCMC) simulations to model asymmetric tail risks (Papamarkou et al., 2022; Martín et al., 2025) and the use of Bayesian non-parametric and semi-parametric expectile models for time-varying ES forecasting (Marcellino and Pfarrhofer, 2024; Chen et al., 2025)—a critical gap remains in practical risk-management implementation. As highlighted by Blue et al. (2024) and Yao (2025), the computational complexity and intense simulation requirements of MCMC and Bayesian model averaging often limit their viability for real-time institutional risk reporting. Similarly, the Basel framework emphasizes transparency and supervisory requirements for market-risk models (Gudiño and Mora, 2022). To address this literature gap, this paper applies daily backtesting data to a computationally efficient, discrete Beta-Binomial conjugate framework for VaR breach-rate calibration.

3. Data and Methodology

3.1. Data Source and Regime Classification

To evaluate the quantitative frameworks under structurally diverse market conditions, all empirical analyses in this study utilize daily closing prices of the S&P 500 index and its primary macro-sector Exchange Traded Funds (ETFs)—namely, Technology (XLK), Financials (XLF),

Healthcare (XLV), and Energy (XLE)—sourced from Yahoo Finance. The historical horizon is symmetrically partitioned into two distinct 5-year regimes to capture potential differences across market regimes: the Pre-Pandemic Regime (January 1, 2014 – December 31, 2018) and the Post-Pandemic Regime (January 1, 2021 – December 31, 2025). The 2019–2020 period is excluded from the out-of-sample evaluation period to facilitate a comparison between relatively stable pre- and post-pandemic regimes rather than to estimate the direct impact of the pandemic shock itself. We acknowledge that excluding this transition period from the evaluation sample constitutes a limitation of the research design.

3.2. The Baseline Portfolio and Tail Risk Formulation

Before implementing empirical backtesting, it is necessary to establish the theoretical baseline for portfolio allocation and tail risk. We consider a financial institution managing a portfolio of n assets. Let w be the $n \times 1$ vector of portfolio weights, μ be the vector of expected returns, and Σ be the $n \times n$ covariance matrix. The optimization problem is formulated to minimize variance subject to a target return m and a full investment constraint:

$$\min_w \frac{1}{2} w' \Sigma w \quad \text{s.t.} \quad w' \mu = m, \quad w' I = 1 \quad (1)$$

Introducing the Lagrangian function with multipliers λ and γ :

$$L(w, \lambda, \gamma) = \frac{1}{2} w' \Sigma w + \lambda(m - w' \mu) + \gamma(1 - w' I) \quad (2)$$

By evaluating the first-order conditions under the informational constants $A = I' \Sigma^{-1} I$, $B = \mu' \Sigma^{-1} I$, and $C = \mu' \Sigma^{-1} \mu$, the unique analytical closed-form solution for the optimal allocation vector w^* is expressed as:

$$w^* = \frac{1}{AC - B^2} \Sigma^{-1} [(A\mu - BI)m + (CI - B\mu)] \quad (3)$$

For ex-post risk measurement, assuming returns follow a parametric distribution $N(\mu, \sigma^2)$, VaR and ES at confidence level c are mathematically expressed as:

$$VaR_c(X) = -[\mu + \sigma \Phi^{-1}(1 - c)] \quad (4)$$

$$ES_c(X) = -[\mu - \sigma \frac{\phi(\Phi^{-1}(c))}{1 - c}] \quad (5)$$

While this analytical baseline provides a deterministic benchmark for asset allocation, it fundamentally assumes a static Gaussian distribution. To address the limitations of these static assumptions in high-volatility regimes, this study transitions to an empirical backtesting framework.

3.3. Volatility Dynamics and Backtesting Framework

To evaluate the performance of the VaR models, we implement a one-day horizon VaR backtesting procedure. The empirical daily returns are projected against the estimated 1-day VaR threshold. We construct and contrast two distinct volatility specifications. The first one is the 21-Day Rolling Window VaR as the Baseline. Representing traditional industry practice, the

baseline standard deviation σ_t is calculated using an unweighted rolling window of the preceding 21 trading days, which is intuitive but structurally suffers from a “ghosting effect” and slow adaptation to sudden market shocks. The benchmark is GARCH(1,1) Conditional VaR to explicitly model volatility clustering. We specify a conditional heteroskedasticity framework where the conditional variance σ_t^2 is updated as:

$$\sigma_t^2 = \omega + \alpha r_{t-1}^2 + \beta \sigma_{t-1}^2 \quad (6)$$

During model estimation, the log-return series is scaled by a factor of 100 to improve numerical convergence of the maximum likelihood function, and the extracted conditional volatility is subsequently rescaled for daily VaR projection. The GARCH(1,1) model assumes a constant conditional mean and Gaussian innovations, with parameters estimated using an expanding window. An initial estimation window of 500 trading days is adopted prior to the evaluation period, followed by daily parameter re-estimation as the window expands. One-step-ahead conditional volatility forecasts are generated for each subsequent trading day and used to construct the 99% Gaussian VaR.

A failure (or breach) occurs when the actual daily loss exceeds the projected VaR. To statistically evaluate the VaR forecasts, we employ two standard backtesting diagnostics. First, the Kupiec Proportion of Failures (POF) test is utilized to determine unconditional coverage—assessing whether the observed violation frequency is statistically consistent with the nominal 1% exception rate. Second, the Christoffersen Markov-chain test is applied to evaluate independence. The Christoffersen test is critical for our framework as it mathematically evaluates whether VaR violations exhibit statistically significant serial dependence or clustering.

3.4. Conjugate Bayesian Calibration

To incorporate the uncertainty surrounding the VaR violation probability into VaR backtesting, we construct a discrete Bayesian updating procedure. The mathematical basis for treating the breach sequence as independent Bernoulli trials is evaluated through the Christoffersen independence test. When the test results do not provide evidence against the independence assumption, we proceed by modeling the daily breach count k over a total window of T days using a Binomial likelihood function governed by an unknown true violation probability θ :

$$f(k | \theta) = \binom{T}{k} \theta^k (1 - \theta)^{T-k} \quad (7)$$

By anchoring the 1% target breach rate in a flexible conjugate Beta prior distribution, $\theta \sim \text{Beta}(\alpha_0, \beta_0)$, the prior density is multiplied by the binomial likelihood to update the system analytically into a posterior Beta distribution:

$$P(\theta | k) \propto \theta^{k+\alpha_0-1} (1 - \theta)^{T-k+\beta_0-1} \sim \text{Beta}(\alpha_0 + k, \beta_0 + T - k) \quad (8)$$

Rather than relying on static frequentist point estimates, this conjugate updating mechanism yields a posterior expected breach rate ($\widehat{\theta}_{post}$). By dividing this posterior mean by the 1% breach rate target ($c = 0.01$), we derive a Bayesian breach-rate adjustment factor to recalibrate internal risk exposures against the uncertainty regarding the model’s exceedance performance (the derivation of the conjugate posterior is provided in Appendix B).

4. Empirical Baseline: Differences in Asset Allocation Across Market Regimes

Before advancing to the dynamic tail-risk backtesting, this section serves as a complementary component to our framework by intuitively illustrating the differences across the pre- and post-pandemic market regimes. We apply the analytical Lagrangian optimization framework presented in Section 3.2 to our universe of four major S&P 500 sector ETFs: Technology (XLK), Financials (XLF), Healthcare (XLV), and Energy (XLE). To ensure cross-regime comparability, the exogenous target portfolio return is uniformly constrained at $m = 10.0\%$, aligning with the historical long-term baseline of the broader U.S. equity market.

4.1. Covariance Structure Across Regimes

To satisfy the prerequisites of deterministic portfolio optimization, we first extract the unconditional covariance matrices (Σ) directly from the daily empirical returns. The structural matrices for the pre-pandemic (2014–2018) and post-pandemic (2021–2025) regimes are presented below:

$$\Sigma_{2014} = \begin{bmatrix} 0.0272 & 0.0195 & 0.0183 & 0.0189 \\ 0.0195 & 0.0276 & 0.0174 & 0.0207 \\ 0.0183 & 0.0174 & 0.0232 & 0.0160 \\ 0.0189 & 0.0207 & 0.0160 & 0.0448 \end{bmatrix} \quad (9)$$

$$\Sigma_{2021} = \begin{bmatrix} 0.0608 & 0.0278 & 0.0181 & 0.0196 \\ 0.0278 & 0.0356 & 0.0163 & 0.0275 \\ 0.0181 & 0.0163 & 0.0210 & 0.0109 \\ 0.0196 & 0.0275 & 0.0109 & 0.0718 \end{bmatrix} \quad (10)$$

A direct comparison of the main diagonals (variances) reveals notable differences across the two subsamples. During the 2014–2018 period, variances were relatively contained. In contrast, the 2021–2025 post-pandemic regime is characterized by a general increase in volatility, with the variance of the Technology sector (XLK) more than doubling from 0.0272 to 0.0608, and Energy (XLE) increasing to 0.0718.

4.2. Optimal Allocation and Parameter Sensitivity

By substituting the empirical parameters into the Lagrangian closed-form solution, the static optimizer resolves the optimal weight vectors (w^*) that minimize overall portfolio risk (σ_{Π}) while strictly satisfying the 10.0% return constraint. The comparative allocation results are summarized in Table 1.

Table 1. Empirical Parameters and Optimal Asset Allocation across Market Regimes

Sector / Metric	Pre-Pandemic (2014–2018)		Post-Pandemic (2021–2025)	
	Expected Return (μ)	Optimal Weight (w^*)	Expected Return (μ)	Optimal Weight (w^*)
Technology (XLK)	12.75%	25.62%	17.31%	-0.21%
Financials (XLF)	7.61%	21.39%	14.62%	8.97%
Healthcare (XLV)	10.26%	50.93%	8.06%	80.76%
Energy (XLE)	-5.75%	2.05%	21.14%	10.48%
Minimum Portfolio Risk (σ_{Π})		14.32%		13.94%

The results in Table 1 provide an intuitive visualization of how static allocations respond to changing parameters. In the pre-pandemic era, the optimizer achieved the baseline target through a balanced structure, allocating 50.93% to Healthcare while assigning only 2.05% to the Energy sector.

However, when the post-pandemic sample moments are introduced, the static optimizer produces a markedly different allocation. To maintain the 10.0% target return without inflating portfolio variance, the algorithm reduces the allocation to Financials, assigns a near-zero weight to Technology (-0.21%), and concentrates more capital in the Healthcare sector (80.76%). Despite this concentration, the theoretical minimum portfolio volatility decreases slightly from 14.32% to 13.94%.

Taken together, the results illustrate the sensitivity of static mean-variance allocations to regime-specific sample moments. The substantial change in optimal weights between the two periods, particularly the pronounced post-pandemic concentration in Healthcare, arises from changes in the estimated expected returns and covariance structure across regimes. This sensitivity highlights an important limitation of interpreting static portfolio weights as stable allocations when the underlying market environment changes. The analysis in Section 5 therefore shifts from static allocation to ex-post tail-risk backtesting and Bayesian calibration using the broader S&P 500 index.

5. Empirical Results: Tail Risk Backtesting and Bayesian Calibration

This section delivers the primary empirical contribution of this paper, shifting the analytical lens from ex-ante sector allocation to ex-post tail-risk backtesting of the broader market. Utilizing daily log-returns of the S&P 500 index, we evaluate the predictive accuracy of the traditional 21-day Rolling Window VaR against the conditional GARCH(1,1) VaR at the 99% confidence level. The backtesting is conducted symmetrically between two 5 years across the pre-pandemic period (01/01/2014–12/31/2018; 1235 trading days) and the post-pandemic period (01/01/2021–12/31/2025; 1232 trading days) using strictly non-overlapping 1-day horizons.

5.1. Volatility Clustering and VaR Backtesting


We first test for the presence of autoregressive conditional heteroskedasticity using the ARCH-LM test to empirically justify the transition from static metrics to conditional volatility models. The test yields a statistic of 196.32 ($P - value = 9.42 \times 10^{-37}$) for the 2014–2018 period and 165.41 ($P - value = 2.48 \times 10^{-30}$) for the 2021–2025 period. These highly significant results strongly reject the null hypothesis of homoskedasticity, providing evidence of conditional heteroskedasticity across both periods and supporting the use of the GARCH(1,1) specification. Subsequently, we subject the realized VaR breach sequences to standard diagnostic tests. The Kupiec Proportion of Failures (POF) test evaluates unconditional coverage (whether the total number of breaches aligns with the 1% expectation), while the Christoffersen test evaluates serial independence (whether consecutive breaches are clustered or independent). The comprehensive backtesting results are summarized in Table 2.

Table 2. S&P 500 Unconditional and Conditional VaR Backtesting Diagnostics

Regime (Total Days)	Risk Model	Breaches (k)	Expected	Kupiec POF Stat	Kupiec P-value	Christoffersen Stat	Christoffersen P-value
Pre-Pandemic (2014–2018, T=1235)	21-Day Rolling VaR	31	12.35	20.0465	0.0000** *	7.3695	0.0066***
	GARCH(1, 1) VaR	25	12.35	10.0923	0.0015** *	21.3751	0.0000***
Post-Pandemic (2021–2025, T=1232)	21-Day Rolling VaR	31	12.32	20.1389	0.0000** *	0.0598	0.8068
	GARCH(1, 1) VaR	23	12.32	7.4502	0.0063** *	0.578	0.4471

Note: *** indicates statistical significance at the 1% level, rejecting the respective null hypotheses. For the Kupiec test, H_0 : Accurate unconditional coverage ($E[k]/T = 0.01$). For the Christoffersen test, H_0 : Serial independence of breach events.

The empirical evidence in Table 2 demonstrates that both frequentist models exhibit more VaR violations than implied by the 1% target rate across the two periods. In both periods, the 21-Day Rolling Window VaR generated 31 exceptions against an expectation of approximately 12.3, yielding Kupiec P -values of 0.0000. While the GARCH(1,1) specification improved the shock absorption capacity by dynamically updating volatility (reducing breaches to 25 and 23, respectively), its Kupiec statistics remain significant at the 1% level ($P = 0.0015$ and $P = 0.0063$). The persistent failure to achieve unconditional coverage is consistent with inadequately captured tail risk. We visually examine these market exceptions under different econometric structures; the historical price trajectories overlaid with the precise 1-day VaR breaches are mapped sequentially in Figure 1 for the pre-pandemic era and Figure 2 for the post-pandemic cycle.






Figure 1. S&P 500 Index Trajectory and 1-Day 99% VaR Breaches (Pre-Pandemic Regime, 2014–2018)




Figure 2. S&P 500 Index Trajectory and 1-Day 99% VaR Breaches (Post-Pandemic Regime, 2021–2025)

The Christoffersen conditional coverage results and the visual layouts in Figures 1 and 2 reveal differences in breach dependence across the two periods. During the 2014–2018 period, the Christoffersen P -values for both models are highly significant ($P = 0.0066$ and $P = 0.0000$), meaning the null hypothesis of independence is strictly rejected. As visually demonstrated in Figure 1, exceptions in this era appear clustered, typically triggered by isolated, massive systemic drawdowns (such as the market correction in early 2018) where the model failed to react in time, leading to consecutive breaches.

However, in the 2021–2025 post-pandemic period, the dynamic shifts fundamentally. The Christoffersen P -values leap to 0.8068 for the Rolling VaR and 0.4471 for the GARCH(1,1) VaR. In the language of statistical inference, both models strictly fail to reject the null hypothesis of serial independence. This is also visually consistent with the spatial distribution of the exception marks in Figure 2. In both panels (Rolling and GARCH), the red cross marks are more dispersed rather than visibly concentrated into dense clusters. Even during and after the severe macroeconomic tightening cycle of 2022, the strict 1-day horizon results do not provide evidence against the independence assumption for the post-pandemic breach sequence. Accordingly, an independent-Bernoulli formulation is used as a working assumption for the subsequent Bayesian Beta-Binomial framework to update the uncertainty surrounding the VaR violation probability. Conversely, for the pre-pandemic period, the Beta-Binomial posterior estimates are retained as descriptive evidence.

5.2. Bayesian Posterior Calibration

Building upon the empirical backtesting diagnostics, we implement the Bayesian Beta-Binomial conjugate framework (derived in Section 3.4) to quantify the uncertainty surrounding the VaR violation probability. To anchor the calibration within existing institutional frameworks, we establish the prior distribution based on the Basel Committee’s standard expectation (Fukui and Basso, 2026). Under the 99% confidence level, the expected exception rate is strictly $c = 0.01$, parameterized as an informative Beta prior, $\theta \sim \text{Beta}(1,99)$. By integrating this 1% target breach rate prior with the realized Binomial likelihood over the respective testing horizon s , the framework continuously updates into a posterior distribution. The complete analytical results, including the posterior parameters, the 95% credible intervals, and the derived Bayesian breach-rate adjustment factors, are consolidated in Table 3.

Table 3. Bayesian Posterior Distributions and Breach-Rate Adjustment Factors

Regime	Risk Specification Model	Empirical Breaches (k)	Posterior Distribution	95% Credible Interval	Posterior Mean ($\hat{\theta}_{\text{post}}$)	Bayesian Breach-Rate Adjustment Factor
Pre-Pandemic (2014–2018)	21-Day Rolling VaR	31	Beta(32,1303)	[1.65%, 3.28%]	2.397%	2.40x
	GARCH(1,1) VaR	25	Beta(26,1309)	[1.25%, 2.76%]	1.947%	1.95x
Post-Pandemic (2021–2025)	21-Day Rolling VaR	31	Beta(32,1300)	[1.65%, 3.29%]	2.402%	2.40x
	GARCH(1,1) VaR	23	Beta(24,1308)	[1.16%, 2.58%]	1.801%	1.80x

Note: The prior distribution is uniformly set to Beta(1,99) across all models. The Bayesian breach-rate adjustment factor is calculated as the ratio of the Posterior Mean to the 1.0% target breach rate. For the pre-pandemic period, the reported posterior estimates are retained as descriptive evidence.

The calibration results in Table 3 reveal differences in the estimated VaR violation probabilities across model specifications and periods. First, the 21-Day Rolling Window VaR exhibits a relatively stable breach-rate adjustment across the two periods. Despite the differences in market conditions between the pre-pandemic and post-pandemic periods, the Rolling VaR generated identical breach counts ($k = 31$) coincidentally, resulting in a nearly identical posterior factor of 2.40x for both periods. This cross-regime consistency indicates that the posterior estimate of the violation probability remains elevated relative to the 1% target under this specification.

Second, the GARCH(1,1) specification produces lower posterior violation probabilities than the Rolling VaR in both periods. By dynamically adjusting to volatility clustering, the GARCH(1,1) specification successfully compressed the posterior mean, lowering the breach-rate adjustment factor to 1.95x in the pre-pandemic period and further down to 1.80x in the post-pandemic period. The difference between the 2.40x and 1.80x factors indicates a lower posterior breach-rate adjustment for the GARCH specification in the post-pandemic period.

Finally, the fact that the GARCH(1,1) posterior mean remains above the 1.0% target breach rate indicates that the conditional volatility specification does not fully eliminate the observed excess in VaR violations relative to the target rate. These results are consistent with residual tail-risk forecasting error and uncertainty surrounding the model's exceedance performance. Accordingly, the Bayesian framework provides an ex-post calibration perspective on observed VaR violations and may complement conventional backtesting in assessing model exceedance performance.

5.3. Robustness Check of the Prior

A frequent critique of Bayesian inferential frameworks is the inherent subjectivity involved in selecting the prior distribution. To assess the sensitivity of the posterior estimates to prior specification, we conduct a sensitivity analysis across three distinct Beta prior specifications for the post-pandemic period (2021–2025, $T = 1232$). We first introduce the baseline regulatory anchored prior, parameterized as $\text{Beta}(1, 99)$, which explicitly reflects the Basel Committee's 1% expected exception rate with a moderate informational weight. To contrast this baseline, we apply an Uninformative Flat prior, $\text{Beta}(1, 1)$, which assumes a uniform probability across all possible exception rates and represents complete initial ignorance. Finally, we incorporate a strong regulatory prior, $\text{Beta}(5, 495)$. This specification strictly maintains the 1.00% expected target but applies an informational weight equivalent to 500 trading days, representing a highly rigid institutional stance that strongly trusts the theoretical expectation over new empirical observations. The interaction between these varied priors and the empirical breach counts is presented in Table 4.

Table 4. Prior Sensitivity Matrix and Posterior Breach-Rate Adjustment Factors (2021–2025)

Prior Specification	Prior Expected Mean	Rolling VaR (k=31) Posterior Mean	Rolling VaR Breach-Rate Adjustment Factor	GARCH(1,1) VaR (k=23) Posterior Mean	GARCH(1,1) Breach-Rate Adjustment Factor
Regulatory Anchored Beta(1,99)	1.00%	2.40%	2.40x	1.80%	1.80x
Uninformative Flat Beta(1,1)	50.00%	2.59%	2.59x	1.94%	1.94x
Strong Regulatory Beta(5,495)	1.00%	2.08%	2.08x	1.62%	1.62x

Note: The Breach-Rate adjustment factor is calculated by dividing the Posterior Mean by the 1.0% target exception rate.

The sensitivity matrix in Table 4 illustrates the sensitivity of the posterior estimates to the prior specification. For the 21-Day Rolling Window VaR, despite the substantial differences in the prior means—ranging from a completely uninformative 50.00% mean to a highly rigid 1.00% strong anchor—the resulting posterior breach-rate adjustment factors fall within the 2.08x to 2.59x range. A parallel range of posterior factors is observed for the GARCH(1,1) specification, more tightly bounded between 1.62x and 1.94x.

This relatively narrow range suggests that the large empirical sample size ($T = 1232$ trading days) generates a dominant binomial likelihood function that mitigates localized subjective biases. Even when subjected to the rigid strong regulatory prior ("Beta" (5,495)), the empirical breach counts raise the posterior mean above the 1.0% target exception rate consistently. This prior-sensitivity analysis therefore indicates that the direction and broad magnitude of the posterior breach-rate adjustment are not driven solely by the baseline prior specification.

5.4. Dynamic Bayesian Calibration

While the static posterior analysis in Section 5.2 establishes an aggregate parameterization for the post-pandemic period (2021–2025), a static framework cannot capture the operational reality of dynamic financial supervision. To bridge this gap, we implement a day-by-day dynamic recursive Bayesian updating procedure. Under this continuous inferential scheme, the posterior probability distribution derived at the conclusion of day t automatically updates to serve as the prior probability distribution for day $t+1$. This recursive architecture allows the Bayesian breach-rate adjustment factor to adjust sequentially as empirical VaR outcomes are observed. Figure 3 plots the daily trajectory of the dynamic Bayesian breach-rate adjustment factors derived from both the 21-Day Rolling Window VaR and the GARCH(1,1) conditional VaR frameworks.




Figure 3. Dynamic Evolution of Bayesian Breach-Rate Adjustment Factors: 21-Day Rolling VaR vs. GARCH(1,1) VaR (2021–2025)

The empirical trajectories illustrated in Figure 3 exhibit a distinct, asymmetric “sawtooth” pattern that characterizes the continuous learning mechanism of the conjugate Beta-Binomial engine. Initiated at the 1.0x baseline corresponding to the 1% target breach probability, each sudden, vertical discontinuity (spike) corresponds to an observed daily VaR breach, where the binary breach indicator variable x_t takes a value of one ($x_t = 1$) to denote a VaR breach. In these instances, the framework increases the posterior estimate of the VaR violation probability. Conversely, during extended intervals without observed VaR breaches, where the model successfully encapsulates empirical returns and the indicator variable remains zero ($x_t = 0$), the curves display a smooth, asymptotic mathematical decay. This structural property reflects the gradual updating of the posterior toward lower estimated violation probabilities as additional non-breach observations accumulate.

A comparative evaluation of the two frameworks in Figure 3 illustrates differences in the dynamic paths of the posterior breach-rate adjustment factors. Because the GARCH(1,1) specification dynamically models the conditional variance, it produces fewer breaches than the 21-Day Rolling Window specification over the post-pandemic evaluation period, resulting in a generally lower adjustment factor for the GARCH framework.

The shaded region between the two trajectories illustrates the difference in the estimated breach-rate adjustment factors between the two specifications. By transitioning from a historical moving average to a conditional model, the GARCH framework produces a lower posterior breach-rate adjustment in this empirical application, converging to approximately 1.80x compared with 2.40x for the Rolling Window specification. Nonetheless, because the GARCH curve remains persistently above the 1.0x baseline across much of the testing horizon, the visual evidence shows that observed VaR violations remain above the 1% target under the present empirical application.

6. Recommendations for Risk Management in the Post-Pandemic Period

The theoretical derivation of portfolio optimization and the subsequent empirical evidence of tail risk calibration provide more than a mere mathematical abstraction; they illustrate practical differences in how risk estimates and portfolio allocations respond to changing market conditions. Rather than assuming stable market parameters across periods, the findings suggest that institutions may benefit from combining static allocation analysis with dynamic risk monitoring and model calibration. The findings of this paper offer three distinct avenues for institutional risk mitigation, directly bridging our empirical diagnostics with practical risk management.

6.1. Rethinking Asset Allocation: The Sensitivity of Static Covariance-Based Allocation

The optimization framework established in Section 4 empirically demonstrates that the optimal asset allocation weight vector w^* is critically sensitive to the covariance matrix Σ . In conventional practice, many institutions utilize historical look-back periods to estimate a static Σ , which may produce materially different allocations when the underlying sample moments change across market periods. Our comparison across the pre-pandemic and post-pandemic periods shows that the static optimizer substantially restructured the portfolio, increasing the Healthcare allocation to 80.76% while assigning a near-zero weight to Technology to satisfy the baseline return constraint.

To mitigate this, modern institutions can consider adopting dynamic regime-switching strategies that utilize dynamic conditional correlation models to update covariance inputs on a near real-time basis. Rather than relying on categorical equity or bond exposures, institutions are encouraged to utilize algorithmic factor-based risk decomposition. Recognizing that cross-asset correlations may change materially during periods of market stress, utilizing adaptive covariance structures may help reduce the risks associated with large and abrupt changes in portfolio exposures when historical allocation relationships become less representative of current market conditions.

6.2. Upgrading Risk Measurement: From Unconditional Historical to Conditional Volatility

The VaR backtesting diagnostics in Section 5 highlight the limitations of relying on standard frequentist risk metrics. The pervasive 21-Day Rolling Window VaR exhibited a higher frequency of VaR violations, failing the Kupiec unconditional coverage test by generating 31 extreme breaches against an expectation of 12.3. This is consistent with inadequately captured tail risk and indicates that unconditional historical volatility measures may not fully reflect changing market conditions. Financial institutions are advised to replace static rolling windows with conditional heteroskedasticity specifications, such as the GARCH(1,1) framework examined in this study. By parameterizing volatility clustering, the conditional model reduced empirical breaches in the present application, although unconditional coverage was not achieved. Furthermore, because even the GARCH model exhibited residual tail exceptions, institutions may consider ES, as theoretically derived in Section 3, as a complementary risk measure alongside conditional VaR. Integrating ES can provide additional information on the magnitude of losses

beyond the VaR threshold and may therefore complement VaR-based monitoring and stress-testing exercises.

6.3. Using Dynamic Bayesian Breach-Rate Calibration for Risk Monitoring

Market risk management may benefit from timely updating when model exceedance performance changes over time. Our findings suggest the potential use of the Beta-Binomial conjugate calibration framework as a transparent, recursively updated risk-monitoring tool. Rather than relying on periodic qualitative reviews, institutions may implement the Bayesian mechanism as an automated mathematical monitoring mechanism. Crucially, the empirical Christoffersen results for the post-pandemic period do not provide evidence against the independence assumption, which is consistent with using the day-by-day recursive updating procedure as a working formulation in this period. When a daily VaR breach materializes, the framework automatically updates the posterior probability, increasing the corresponding breach-rate adjustment factor, as visualized by the asymmetric “sawtooth” pattern in Section 5.4. For instance, the empirical testing yielded a final GARCH breach-rate adjustment factor of approximately 1.80x, compared with 2.40x for the 21-Day Rolling Window specification.

Furthermore, as indicated by the prior sensitivity matrix in Section 5.3, this calibration varies within a relatively limited range under the three prior specifications considered, stabilizing between 1.62x and 1.94x for GARCH. Conversely, during extended periods without observed breaches, the mathematical decay of the posterior mean reduces the corresponding breach-rate adjustment factor. Taken together, these results suggest that the proposed Bayesian mechanism may provide a transparent ex-post calibration and monitoring tool for observed VaR exceedance performance.

7. Conclusion

This paper has synthesized a multi-layered framework for risk management, bridging the historical gap between qualitative governance and the quantitative imperatives of modern financial mathematics. By integrating classical optimization techniques with tail-risk measures and Bayesian inference, we have provided a framework for examining risk management challenges under post-pandemic market conditions. Our application of Lagrangian multipliers demonstrates that portfolio optimization remains a mathematically deterministic process; however, its institutional utility is sensitive to the estimated means and covariance structure underlying the allocation problem.

The theoretical treatment of Expected Shortfall (ES), together with the empirical VaR backtesting of S&P 500 returns, highlights the limitations of Gaussian-based models in capturing observed tail-risk behavior. The severe “recognition lag” observed in the 21-Day Rolling Window VaR—generating 31 extreme breaches against an expectation of 12.3—is consistent with inadequately captured tail risk, while the ARCH-LM results provide evidence of conditional heteroskedasticity and volatility clustering. While conditional heteroskedasticity models like GARCH(1,1) reduced the number of observed breaches, their residual tail exceptions indicate that unconditional coverage was not achieved.

Consequently, the core contribution of this paper lies in the Bayesian recalibration of model reliability. By moving beyond frequentist backtesting to a Beta-Binomial conjugate prior framework, we characterize uncertainty surrounding the VaR violation probability. The Bayesian breach-rate adjustment factors, converging to 1.80x for the GARCH model and 2.40x for the 21-Day Rolling Window specification, indicate elevated posterior breach-rate estimates relative to the 1% target in the present empirical application. This suggests a potential role for dynamic ex-post calibration alongside conventional model monitoring rather than relying solely on static model estimates.

Ultimately, the findings point to three complementary directions for risk management: rigorous optimization utilizing coherent risk metrics like ES to account for tail severity, conditional volatility modeling and empirical backtesting to assess exceedance performance, and automated Bayesian updating to incorporate new VaR outcomes into posterior breach-probability estimates. As financial markets continue to evolve toward algorithmic and geopolitical complexity, these complementary components may provide a mathematically grounded approach to institutional risk monitoring in post-pandemic markets.

Declaration of Research and Publication Ethics

This study which does not require ethics committee approval and/or legal/specific permission complies with the research and publication ethics.

Researcher's Contribution Rate Statement

H.J., development of the concept of the article, theoretical part, conceptualization, collection of statistical data, writing – original draft, collection and analysis of literature, writing–review & editing, description of results, formation of tables, visualization.

H.S., theoretical part, econometric analysis, review, description of results.

S.A., theoretical part, econometric analysis, collection and analysis of literature, writing–review & editing.

J.H., writing–review & editing, description of results, formation of tables, visualization.

Declaration of Researcher's Conflict of Interest

There is no potential conflict of interest in this study.

Declaration of Artificial Intelligence Usage

Artificial intelligence tools were used in the language editing of this manuscript. Apart from this, the authors did not use any artificial intelligence tools during the preparation of this manuscript.

References

- Acerbi, C. and Tasche, D. (2002). Expected shortfall: A natural coherent alternative to value at risk. *Economic Notes*, 31(2), 379-388. <https://doi.org/10.1111/1468-0300.00091>
- Arfaoui, N. and Yousaf, I. (2022). Impact of COVID-19 on volatility spillovers across international markets: Evidence from VAR asymmetric BEKK GARCH model. *Annals of Financial Economics*, 17(01), 2250004. <https://doi.org/10.1142/S201049522250004X>
- Artzner, P., Delbaen, F., Eber, J.M. and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228. <https://doi.org/10.1111/1467-9965.00068>

- Babic, B. (2019). A theory of epistemic risk. *Philosophy of Science*, 86(3), 522-550. <https://doi.org/10.1086/703552>
- Balakrishnan, K. and Ertan, A. (2021). Credit information sharing and loan loss recognition. *The Accounting Review*, 96(4), 27-50. <https://doi.org/10.2308/tar-2017-0244>
- Beckmann, J., Belke, A. and Dubova, I. (2022). What drives updates of inflation expectations? A Bayesian VAR analysis for the G-7 countries. *The World Economy*, 45(9), 2748-2765. <https://doi.org/10.1111/twec.13241>
- Blue, G., Faraji, O., Khotanlou, M. and Rezaee, Z. (2024). A corporate risk assessment and reporting model in emerging economies. *Journal of Applied Accounting Research*, 25(4), 783-811. <https://doi.org/10.1108/jaar-02-2023-0047>
- Carriero, A., Clark, T.E. and Marcellino, M. (2024). Capturing macro-economic tail risks with Bayesian vector autoregressions. *Journal of Money, Credit and Banking*, 56(5), 1099-1127. <https://doi.org/10.1111/jmcb.13121>
- Chen, C.W.S, Chen, P.H. and Hsu, Y.L. (2025). Bayesian forecasting of value-at-risk and expected shortfall in cryptocurrency markets: A nonlinear semi-parametric framework. *Applied Stochastic Models in Business and Industry*, 41(1), 2926. <https://doi.org/10.1002/asmb.2926>
- Dimitriadis, K.A., Koursaros, D. and Savva, C.S. (2025). The influential impacts of international dynamic spillovers in forming investor preferences: A quantile-VAR and GDCC-GARCH perspective. *Applied Economics*, 57(45), 7175-7195. <https://doi.org/10.1080/00036846.2024.2387868>
- Fabozzi, F.J., Markowitz, H.M. and Gupta, F. (2008). Portfolio selection. *Handbook of Finance*, 2, 3-13. <https://doi.org/10.1002/9780470404324.hof002001>
- Feldkircher, M., Huber, F., Koop, G. and Pfarrhofer, M. (2022). Approximate Bayesian inference and forecasting in huge-dimensional multicountry VARs. *International Economic Review*, 63(4), 1625-1658. <https://doi.org/10.1111/iere.12577>
- Fukui, M. and Basso, L.F.C. (2026). The relevance of expected shortfall models in different time window sizes. *International Journal of Financial Studies*, 14(2), 42. <https://doi.org/10.3390/ijfs14020042>
- Gudiño, J.J.C. and Mora, J.A.N. (2022). Machine learning models, risk management current regulation and perspectives. In J.A.N. Mora & M.B.M. Aragón (Eds.), *Data analytics applications in emerging markets* (pp. 49-73). Singapore: Springer Nature. https://doi.org/10.1007/978-981-19-4695-0_3
- Huang, Y., Wang, H., Chen, Z., Feng, C., Zhu, K., Yang, X. and Yang, W. (2024). Evaluating cryptocurrency market risk on the blockchain: An empirical study using the ARMA-GARCH-VaR model. *IEEE Open Journal of the Computer Society*, 5, 83-94. <https://doi.org/10.1109/ojcs.2024.3370603>
- Jui, F.N., Hossain, M.J., Das, A., Sultana, N. and Islam, M.K. (2024). Analyzing the impact of remittance, FDI and inflation rate on GDP: A comparative study of Bangladesh, Pakistan and Sri-Lanka using VAR and BEKK-GARCH approach. *Heliyon*, 10(11), e31381. <https://doi.org/10.1016/j.heliyon.2024.e31381>
- Kakade, K., Jain, I. and Mishra, A.K. (2022). Value-at-risk forecasting: A hybrid ensemble learning GARCH-LSTM based approach. *Resources Policy*, 78, 102903. <https://doi.org/10.1016/j.resourpol.2022.102903>
- Kan, R. and Lassance, N. (2025). Optimal portfolio choice with fat tails and parameter uncertainty. *Journal of Financial and Quantitative Analysis*, 60(8), 3753-3790. <https://doi.org/10.1017/s0022109025101622>
- Kuang, W. (2024). High-frequency enhanced VaR: A robust univariate realized volatility model for diverse portfolios and market conditions. *PLOS one*, 19(5), 0303962. <https://doi.org/10.1371/journal.pone.0303962>
- Lam, M.D. (2025). Estimation and backtesting of financial risk measures in the Vietnamese stock market: The debate between value-at-risk and expected shortfall. *Global Business and Economics Review*, 33(3-4), 423-443. <https://doi.org/10.1504/gber.2025.148351>

- Le, T.H. (2024). Forecasting VaR and ES in emerging markets: The role of time-varying higher moments. *Journal of Forecasting*, 43(2), 402-414. <https://doi.org/10.1002/for.3039>
- Maharana, N., Panigrahi, A.K. and Chaudhury, S.K. (2024). Volatility persistence and spillover effects of Indian market in the global economy: A pre-and post-pandemic analysis using VAR-BEKK-GARCH model. *Journal of Risk and Financial Management*, 17(7), 294. <https://doi.org/10.3390/jrfm17070294>
- Maharana, N., Panigrahi, A.K., Chaudhury, S.K., Uprety, M., Barik, P. and Kulkarni, P. (2025). Economic resilience in post-pandemic India: Analysing stock volatility and global links using VAR-DCC-GARCH and wavelet approach. *Journal of Risk and Financial Management*, 18(1), 18. <https://doi.org/10.3390/jrfm18010018>
- Mallela, J. and Gupta, M. (2024). Value at risk performance and backtesting of cryptocurrencies. *Management Quest*, 6, 51-73. Retrieved from <https://www.dsims.org.in/>
- Marcellino, M. and Pfarrhofer, M. (2024). Bayesian nonparametric methods for macroeconomic forecasting. In M.P. Clements & A.B. Galvão (Eds.), *Handbook of research methods and applications in macroeconomic forecasting* (pp. 90-125). Edward Elgar Publishing. <https://doi.org/10.4337/9781035310050.00009>
- Martín, J., Parra, M.I., Pizarro, M.M. and Sanjuán, E.L. (2025). A new Bayesian method for estimation of value at risk and conditional value at risk. *Empirical Economics*, 68(3), 1171-1189. <https://doi.org/10.1007/s00181-024-02664-2>
- Mufaridho, L.M.W., SH, K.N., Jalaludin, P. and Pratama, M.I. (2026). Backtesting of the Value-at-Risk based on GARCH model (VaR-GARCH) in measuring stock market risk. *Journal of Mathematics, Computations and Statistics*, 9(1), 39-47. <https://doi.org/10.35580/eyqd4722>
- Muhtarullo, F., Humairo, R.I.S. and Meiza, A. (2026). Option pricing for exchange rate hedging: Evaluation of value at risk, sharpe ratio, and backtesting with Kupiec and Christoffersen tests. *JINAV: Journal of Information and Visualization*, 7(1), 1-10. <https://doi.org/10.35877/454RI.jinav4667>
- Naifar, N. (2025). Mapping systemic tail risk in crypto markets: Defi, stablecoins, and infrastructure tokens. *Journal of Risk and Financial Management*, 18(6), 329. <https://doi.org/10.3390/jrfm18060329>
- Nasrudin, M., Setyowati, E. and Wara, S.S.M. (2024). Application of VAR-GARCH for modeling the causal relationship of stock prices in the mining sub-sector. *Jurnal Varian*, 8(1), 89-96. <https://doi.org/10.30812/varian.v8i1.4239>
- Papamarkou, T., Hinkle, J., Young, M.T. and Womble, D. (2022). Challenges in Markov chain Monte Carlo for Bayesian neural networks. *Statistical Science*, 37(3), 425-442. <https://doi.org/10.1214/21-sts840>
- Yao, Y. (2025). Applications of Bayesian inference in financial econometrics: A review. *Economics and Management Innovation*, 2(2), 45-52. <https://doi.org/10.71222/mz71ts21>
- Zhu, Y., Taasim, S.I. and Daud, A. (2025). Volatility modeling and tail risk estimation of financial assets: Evidence from gold, oil, Bitcoin, and stocks for selected markets. *Risks*, 13(7), 138. <https://doi.org/10.3390/risks13070138>

APPENDIX

A. Expected Shortfall (ES) Mathematics Proof

The original equations defining the Expected Shortfall and Value-at-Risk under a Normal distribution are given by:

$$ES_c(X) = E[X|X \leq VaR_c(X)] \text{ and } VaR(X) = \mu + \sigma\Phi^{-1}(1 - c)$$

The following is the step-by-step rigorous mathematical derivation of the closed-form ES solution. By definition, ES is the integral of the VaR function over the tail probabilities:

$$\begin{aligned} ES_c(X) &= \frac{1}{1-c} \int_c^1 VaR_u(X) du \\ &= \frac{1}{1-c} \int_c^1 (\mu + \sigma \Phi^{-1}(1-u)) du \\ &= \mu + \frac{\sigma}{1-c} \int_c^1 \Phi^{-1}(1-u) du \end{aligned}$$

To solve the integral $\int_c^1 \Phi^{-1}(1-u) du$, we apply integration by substitution. Let $u = \Phi(z)$, which implies the differential is $du = \phi(z) dz$. Consequently, the integration limits must be transformed: the lower limit $u = c$ becomes $z = \Phi^{-1}(c)$, and the upper limit $u = 1$ becomes $z = \infty$. Furthermore, utilizing the symmetry of the standard normal distribution, we know that $\Phi^{-1}(1 - \Phi(z)) = -z$. Substituting these components back into the integral yields:

$$\begin{aligned} \int_c^1 \Phi^{-1}(1-u) du &= \int_{\Phi^{-1}(c)}^{\infty} (-z) \phi(z) dz \\ &= \int_{\Phi^{-1}(c)}^{\infty} -z \left(\frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} \right) dz \end{aligned}$$

Using the fundamental calculus rule $\int z e^{-\frac{z^2}{2}} dz = -e^{-\frac{z^2}{2}}$, we evaluate the definite integral:

$$\begin{aligned} \int_{\Phi^{-1}(c)}^{\infty} -z \left(\frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} \right) dz &= \left[\frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} \right]_{\Phi^{-1}(c)}^{\infty} \\ &= 0 - \phi(\Phi^{-1}(c)) \\ &= -\phi(\Phi^{-1}(c)) \end{aligned}$$

Substituting this integrated tail result back into the main ES equation provides the final closed-form analytical solution:

$$ES_c(X) = \mu - \sigma \frac{\phi(\Phi^{-1}(c))}{1-c}$$

B. Bayesian Framework and Posterior Derivation

This appendix details the statistical formulation used in Section 5 to quantify the uncertainty surrounding the VaR violation probability. We employ a Conjugate Prior approach to derive the posterior distribution of the true underlying breach probability, θ . The mathematical demonstration below uses the empirical results of the GARCH(1,1) model during the Post-Pandemic Regime as the illustrative example.

B.1. The Likelihood Function

For the post-pandemic example considered here, we model the daily VaR breach event as an independent Bernoulli trial, consistent with the Christoffersen result that the test does not provide evidence against the independence assumption in this period, where a breach occurs with

probability θ and does not occur with probability $1 - \theta$. The total number of breaches k over T observations follows a Binomial distribution. The likelihood function $L(Data | \theta)$ is given by:

$$L(k | T, \theta) = \binom{T}{k} \theta^k (1 - \theta)^{T-k}$$

where $T = 1232$ represents the total trading days in the post-pandemic sample, and $k = 23$ represents the observed empirical breaches generated by the GARCH(1,1) specification.

B.2. The Prior Distribution

To incorporate the 1% target breach rate, we select the Beta distribution as the prior for the unknown parameter θ , as it is the mathematical conjugate prior for the Binomial likelihood:

$$P(\theta | \alpha, \beta) = \frac{\theta^{\alpha-1} (1 - \theta)^{\beta-1}}{B(\alpha, \beta)}$$

We calibrate the hyperparameters α and β to reflect the theoretical Basel expectation. For a 99% confidence level VaR, the expected breach rate is 1.0%. Thus, we anchor the prior expectation $E[\theta] = \frac{\alpha}{\alpha + \beta} = 0.01$. We utilize the baseline prior values of $\alpha_{prior} = 1$ and $\beta_{prior} = 99$, representing a standardized prior belief equivalent to observing 1 breach in 100 theoretical trading days.

B.3. Posterior Derivation

By Bayes' Theorem, the posterior distribution $P(\theta | Data)$ is proportional to the product of the likelihood and the prior:

$$\begin{aligned} P(\theta | Data) &\propto L(Data | \theta) \times P(\theta) \\ &\propto [\theta^k (1 - \theta)^{T-k}] \times [\theta^{\alpha-1} (1 - \theta)^{\beta-1}] \\ &\propto \theta^{(\alpha+k)-1} (1 - \theta)^{(\beta+T-k)-1} \end{aligned}$$

This functional form perfectly aligns with the kernel of a new Beta distribution. Therefore, the posterior distribution is analytically derived as:

$$\theta | Data \sim \text{Beta}(\alpha_{post}, \beta_{post})$$

where the updated parameters are $\alpha_{post} = \alpha_{prior} + k$ and $\beta_{post} = \beta_{prior} + T - k$. Substituting the empirical GARCH(1,1) values:

$$\begin{aligned} \alpha_{post} &= 1 + 23 = 24 \\ \beta_{post} &= 99 + 1232 - 23 = 1308 \end{aligned}$$

It is important to note that, applying this identical mechanism to the 21-Day Rolling VaR framework yields $k = 31$, generating a posterior distribution of Beta(32,1300).

B.4. Posterior Estimators and Breach-Rate Adjustment

The Posterior Mean, utilized as the posterior point estimate of the VaR violation probability, is calculated as:

$$E[\theta | Data] = \frac{\alpha_{post}}{\alpha_{post} + \beta_{post}} = \frac{24}{24 + 1308} \approx 1.801\%$$

The Bayesian Breach-Rate Adjustment Factor (M) is defined as the ratio of this posterior mean to the 1% target breach rate ($\theta_{target} = 0.01$):

$$M = \frac{E[\theta | Data]}{\theta_{target}} = \frac{0.01801}{0.01} \approx 1.80x$$

This final factor provides a mathematically derived summary of the posterior VaR violation probability relative to the 1% target.