Proceedings of the 10th International Conference on Economic Management and Green Development (ICEMGD 2026)
DOI: 10.54254/2754-1169/2026.33934

A Comparative Study of GBM and GARCH Models for Pricing Automatically Redeemable Structured Products—Taking HSBC Trigger Autocallable Notes as an Example

Silong Lin

Lee Shau Kee School of Business and Administration, Hong Kong Metropolitan University, Hong Kong, China
s1336620@live.hkmu.edu.hk

Abstract. Automatically redeemable structured products exhibit path dependence, and their pricing is typically achieved using Monte Carlo simulations, assuming the underlying asset price follows a geometric Brownian motion (GBM); this means that volatility is constant. However, real financial markets exhibit volatility clustering and fat tails, and the constant volatility assumption can lead to pricing biases. This paper takes a trigger autocallable note issued by HSBC and linked to the S&P 500 index as an example. It uses both GBM and GARCH(1,1) models to generate the underlying asset price path, calculates the product's theoretical value and expected loss (ES) using Monte Carlo simulations within a risk-neutral framework, and compares the results from four dimensions: fair value, risk indicators, return distribution, and sample path. The results show that the constant volatility of GBM leads to overly dispersed paths and an overestimation of loss frequency, while GARCH, by characterizing time-varying volatility and mean reversion, provides a risk-return profile that better reflects market realities. Therefore, this paper recommends using the GBM model when the market is stable, or the product structure is simple, and using the GARCH model when the market is volatile, or the product exhibits strong path dependence.

Keywords: Automatically redeemable structured products, GBM model, GARCH model, Monte Carlo simulation, Expected loss

1. Introduction

Monte Carlo simulation is widely used in the pricing research of automatically redeemable structured products. Traditional pricing practices usually assume that the price of the underlying asset follows a geometric Brownian motion, i.e., the volatility is constant. Since the pioneering work of Black and Scholes [1], the GBM assumption has become the benchmark framework for derivative pricing due to its simplicity and analytical tractability. However, a large number of empirical studies have shown that financial asset return series generally exhibit characteristics such as volatility clustering, conditional heteroscedasticity, and fat-tailed distribution, which deviate significantly from the constant volatility assumption of GBM. To characterize the time-varying characteristics of volatility, Engle proposed the autoregressive conditional heteroscedasticity ARCH model [2], which

DOI: 10.54254/2754-1169/2026.33934

was subsequently extended by Bollerslev to the generalized autoregressive conditional heteroscedasticity GARCH model [3]. In the field of option pricing, Mozumder et al. proposed an approximate closed-end option pricing model based on a nonlinear GARCH process, calibrated using S&P 500 index option samples, and found that the model significantly outperformed existing models both in-sample and out-of-sample [4]. In terms of model comparison, Pang and Zhao used S&P 500 index data to comprehensively compare GARCH and autoregressive stochastic volatility (ARSV) models and found that the two models performed very similarly in call option pricing under risk-neutral measures [5]. Sidhu and Saxena used real-time market data from 2024 to compare GBM, GARCH, Heston and Merton jump diffusion models through Monte Carlo simulation and found that the Heston model was closest to the market price, while the GARCH model provided improved volatility prediction [6]. Hansen and Tong combined the time-varying volatility risk-averse pricing kernel with the Heston-Nandi GARCH model to obtain the easy-to-use option pricing model DHNG, which can significantly reduce VIX and option pricing errors [7]. In addition, Wang and Gerlach proposed a Bayesian realization threshold measurement GARCH framework to predict 1% and 2.5% value at risk and expected loss, and found that the framework is better than the traditional GARCH model in tail risk prediction [8].

In the field of structured product pricing, Cui, Li and Zhang proposed a unified pricing framework based on continuous-time Markov chain approximation, which applies to various structured products such as automatic redemption products, and examined dynamic hedging strategies in the presence of transaction costs [9]. Cheng used the GARCH model to evaluate the impact of artificial intelligence-driven market volatility on S&P 500 index option pricing and found that the GARCH model significantly improved the pricing accuracy of short-term call options [10].

While the superiority of the GARCH model in option pricing has been partially supported by empirical evidence, research comparing the GBM and GARCH models in the specific scenario of automatically redeemable structured products remains relatively limited. Furthermore, existing literature lacks an in-depth discussion on the differences in the performance of expected loss as a risk measure under the two models.

Therefore, this paper takes HSBC's trigger autocallable notes linked to the S&P 500 index as an example. It uses both the GBM and GARCH(1,1) models to generate the underlying asset price path, calculates the theoretical value and expected loss of the product through Monte Carlo simulation within a risk-neutral framework, and compares them from four dimensions: fair value, risk indicators, return distribution, and sample path. The aim is to reveal the impact of time-varying volatility assumptions on the pricing of automatically redeemable structured products and provide a reference for model selection in practice.

2. Method and data

2.1. Data sources and explanations

The underlying asset data in this article is the S&P 500 Index. Daily closing levels of the index were obtained from Yahoo Finance using the yfinance library in Python. The sample period is from January 1, 2021 to February 21, 2024, covering trading days for the three years prior to the product pricing date (February 22, 2024), and includes sufficient observations to ensure the reliability of the estimates. The data includes dates and corresponding closing prices. Based on these prices, daily logarithmic returns can be calculated.

The product analyzed in this article is the Trigger Autocallable Notes issued by HSBC USA in February 2024. Its full terms are disclosed in the free-writing prospectus filed with the U.S.

DOI: 10.54254/2754-1169/2026.33934

Securities and Exchange Commission (SEC). For pricing purposes, the following key features of the contract are extracted and summarized in Table 1.

Table 1. Summary of product terms

Underlying Asset	S&P 500 Index
Pricing date	2024/2/22
Settlement date	2024/2/27
Expiry date	2026/2/26
Observation Day	2025/2/28;2025/5/22;2025/8/22;2025/11/24;2026/2/23
Redemption return rate	8.90%,11.125%,13.35%,15.575%,17.80%
Downlink threshold	75.00% of the initial level
initial level	5,096.27(Pricing Day Closing Level)
Principal amount of each note	10USD

2.2. Indicator selection and explanation

To implement the Monte Carlo simulation, this paper defines the following variables and parameters. Price Process of the Indicator: The future path of the S&P 500 index is simulated using different stochastic processes (GBM and GARCH). Let s_t be the index level at time t (in trading days).

The initial level of the note, s_0, is 5,096.27 USD. Assuming there are 252 trading days in a year, the calendar date is converted to the trading day number from the pricing date. Let T_1 , T_2 , ……, T_5 be the trading day numbers for the corresponding 5 observation days.

The redemption returns are denoted as C_1 , C_2 , ……, C_5 with values of 0.0890, 0.11125, 0.1335, 0.15575, and 0.1780, respectively. If the note is redeemed on the corresponding observation date, the corresponding return applies.

The downside threshold D for the note is $0.75S_0$. To facilitate principal calculation, this paper standardizes the principal to 100 in all calculations, and the final return is expressed as per 100 yuan of principal.

2.3. Method introduction

This paper first uses a specified stochastic process to generate a large number (10,000 in this paper) of sample paths for the S&P 500 index over the product's life. Then, for each simulated path, its yield is determined according to the product's contractual rules. Finally, the average yield across all paths is calculated, and this average is the estimated fair value of the note. To focus primarily on model comparisons, discounting is not considered. The GBM model formulas used in this paper are shown in formulas 1, 2, and 3.

$$dS_t=rS_tdt+\sigma S_tdW_t \tag{1}$$

Where μ is the drift rate (expected return), σ is the constant volatility, W_t is the standard Wiener process. In discrete time, with Δt =1/252, can obtain that.

DOI: 10.54254/2754-1169/2026.33934

$$S_{t+\Delta t}=S_t \exp\left(\left(r-\frac{1}{2}\sigma^2\right)\Delta t+\sigma\sqrt{\Delta t}\,\varepsilon_{t+1}\right),\ \varepsilon_{t+1}\sim N(0,1) \tag{2}$$

The parameters μ and σ can be estimated using historical data of the S&P 500-day logarithmic returns over the sample period.

$$\widehat{\sigma}=\sqrt{\frac{1}{T-1}\sum_{t=1}^{T}\left(r_t-\overline{r}\right)^2\times 252} \tag{3}$$

Where $\overline{r}$ is the mean of the daily logarithmic returns.

The GARCH model formulas used in this paper are shown in formulas 4, 5, 6, and 7. In order to capture the time-varying volatility characteristics observed in the financial market, the GARCH(1,1) model is adopted. The model for the logarithmic return rt is set as follows.

$$r_t=r+\sigma_t z_t,\ z_t\sim N(0,1) \tag{4}$$

$$\sigma_t^2=\omega+\alpha(r_{t-1}-r)^2+\beta\sigma_{t-1}^2 \tag{5}$$

Where σ_t^2 represents the conditional variance at time t, and $\omega>0$, $\alpha\geq 0$, $\beta\geq 0$ satisfy $\alpha+\beta<1$ to ensure the stationarity of the volatility process.

Specifically,

$$\sigma_{t+1}^2=\omega+\alpha(\sigma_t z_t)^2+\beta\sigma_t^2 \tag{6}$$

$$r_{t+1}=r+\sigma_{t+1}z_{t+1},\ z_{t+1}\sim N(0,1) \tag{7}$$

The subsequent price path can be derived from formula 8.

$$S_{t+1}=S_t\exp(r_{t+1}) \tag{8}$$

The return calculation method is as follows: for each simulated price path $\{s_0, s_1, s_2, \ldots\ldots, s_T\}$ (where T is the total number of trading days from the maturity date), automatic redemption is checked. If s_{ti} is greater than or equal to s_0, the note is redeemed on that observation date, and the investment return is

DOI: 10.54254/2754-1169/2026.33934

$$\text{Return}=100\times(1+c_i) \tag{9}$$

If not redeemed, on the final observation date T_5, compare the ending s_{t5} with the downward threshold $D=0.75\times S_0$.

If $s_{t5}\geq D$, the investor recovers the entire principal of $100. If St5< D , the investor suffers a principal loss proportional to the index decline, with the return expressed by formula 10.

$$\text{Return}=100\times\frac{S_{\tau_5}}{S_0} \tag{10}$$

To assess downside risk, this paper uses Expected Shortfall (ES) as the primary risk metric. Compared to the traditional Value at Risk (VaR), ES more comprehensively reflects tail risk, especially when the probability of loss is low. ES considers the average level of all losses exceeding the VaR threshold, thus overcoming VaR's deficiency in characterizing tail risk. The 95% expected loss is defined as the average loss amount under the worst-case 5% scenario. Let $P_{0.05}$ be the 5th percentile of the return, then the formula for calculating the 95% ES is shown in formula 11.

$$ES_{95\backslash\%}=\frac{1}{}\sum_{P_i<P_{(0.05)}}(100\text{-}P_i) \tag{11}$$

Where 100 represents the standardized principal amount, 100 - P_i is the loss amount (when $P_i<$ 100). If the return corresponding to the 5th percentile is exactly equal to the principal ($P_{0.05}$ = 100), then ES only calculates the average loss of loss paths strictly below the principal. This indicator is expressed in US dollars, and a larger value indicates a higher tail risk.

This paper compares the GBM and GARCH models from four dimensions. First, it estimates the fair value of notes under different models. Second, it compares the 95% ES (Expected Estimate) of notes under the two models, as well as the probabilities of automatic redemption, principal protection, and loss scenarios. Third, it visually compares the distribution patterns of the paths simulated by the two models through histograms and sample paths. Fourth, comparing the sample path of two models.

This comparison reveals how considering time-varying volatility affects the valuation and risk assessment of automatically redeemable notes. Analyzing the differences between the two models helps assess whether the constant volatility assumption leads to systematic bias and provides a basis for model selection under different scenarios.

3. Result and discussion

Based on the daily logarithmic returns of the S&P 500 from January 1, 2021, to February 21, 2024, this paper estimates the parameters of the GBM and GARCH(1,1) models. Under the risk-neutral pricing framework, the drift term is set to zero, and only the volatility parameter is retained.

DOI: 10.54254/2754-1169/2026.33934

3.1. Monte carlo simulation results

Using the above parameters, 10,000 S&P 500 price paths were generated using the GBM and GARCH models, respectively, with a period from the pricing date of February 22, 2024, to the maturity date of February 26, 2026. The final return for each path was calculated according to the product terms (principal standardized to 100). Key statistical indicators are summarized in Table 2.

Table 2. Comparison of pricing and risk indicators under the two models

Index	GBM	GARCH
Mean Payoff	100.24	102.04
95%ES (loss)	46.05	44.73
Autocall probability	63.0%	67.58%
Breakeven probability	19.8%	19.1%
Loss probability	17.2%	13.4%

3.1.1. Fair value and scenario probability

As shown in Table 2, under a risk-neutral measure, the average return of the GARCH model is $102.04, higher than GBM's $100.24, indicating that the GARCH path generally tends to yield higher returns. In terms of scenario probability, GARCH's automatic redemption probability is 67.5%, higher than GBM's 63.0%; the probability of principal preservation is similar for both (19.1% vs. 19.8%); while the probability of loss is 13.4% for GARCH, lower than GBM's 17.2%. This suggests that the path generated by the GARCH model increases the opportunity for automatic redemption while reducing the possibility of principal loss.

3.1.2. 95% expected loss (ES)

In terms of risk measurement, GBM's 95% expected loss is 46.05, meaning an average loss of $46.05 in the worst-case 5% scenario; GARCH's ES is 44.73, slightly lower than GBM. Although GARCH has a lower probability of loss, its ES is close to GBM's, indicating that when a loss occurs on the GARCH path, the magnitude of the loss is comparable to or even slightly smaller than GBM's. This reveals the difference between the two in terms of tail risk: GBM's risk is mainly manifested in a higher frequency of losses, while GARCH maintains a similar severity of losses while reducing the frequency of losses.

3.1.3. Profit distribution and sample path

Figure 1 shows the histograms of profits from 10,000 simulations under the two models. The profit distribution of GBM has a distinct left tail (corresponding to loss), with the peak occurring at 100 (break-even) and slightly above 100 in the redemption range; the distribution of GARCH also has a left tail, but it is relatively shorter, with a higher peak concentrated above 100, reflecting its higher automatic redemption probability and lower loss frequency.

DOI: 10.54254/2754-1169/2026.33934

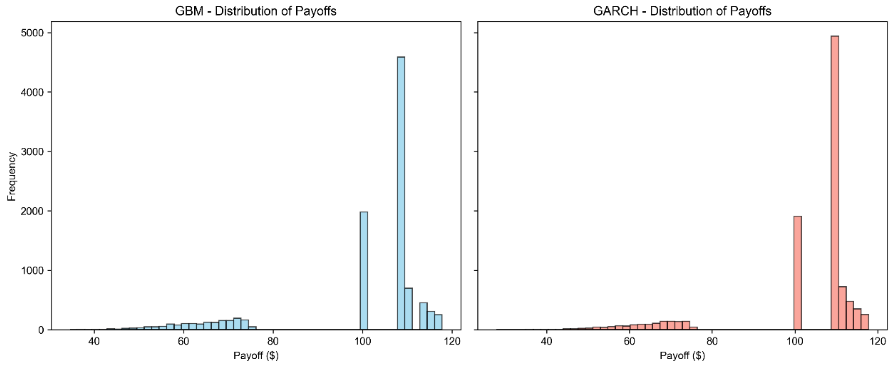

Figure 1. Histograms of simulated returns under GBM and GARCH models

Figure 2 shows a randomly selected GBM path and a GARCH path. The GBM path exhibits relatively uniform volatility, displaying a random walk characteristic; the GARCH path, on the other hand, shows volatility clustering—volatility amplifies and prices fluctuate significantly during certain periods, while remaining relatively stable during other periods. This time-varying volatility characteristic makes the GARCH path more likely to trigger redemptions on certain observation days, but it may also lead to sharp declines in a few cases.

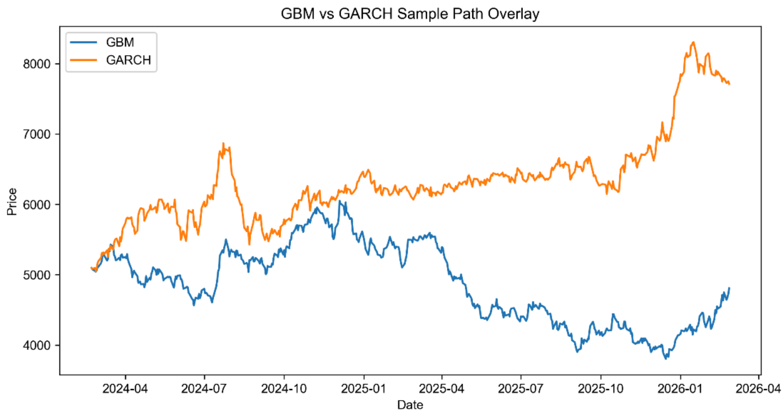

Figure 2. A price path randomly selected from the two models

3.2. Discussion

This study compares the application of GBM and GARCH models in the pricing of automatically redeemable notes, revealing significant differences between the two models within a risk-neutral framework. The different volatility settings of the two models result in these differences.

GBM's constant volatility causes price paths to spread randomly around the initial level without "memory." This setting results in a high degree of path dispersion, also a considerable proportion of paths causing significant declines and losses. Therefore, GBM exhibits a moderate probability of automatic redemption (63.0%) and a relatively high probability of loss (17.2%), with a large ES (46.05). For GARCH, it exhibits a mean-reverting character, when volatility increases, it gradually reverts to its long-term mean. So it is difficult for the path to deviate significantly from the initial level. Consequently, the probability of automatic redemption is slightly higher (67.5%), and the probability of loss is lower (13.4%). However, occasional volatility spikes can still lead to severe

DOI: 10.54254/2754-1169/2026.33934

declines in a few paths. These rare but dramatic losses were captured by ES at 44.73, indicating that GARCH reduces the frequency of losses, but the severity of those losses is comparable to GBM.

Based on the above comparison, when the market is in a stable period, volatility does not show obvious clustering, or when a quick and rough estimate of product value is needed, GBM is sufficient due to its simplicity and computational efficiency. However, it should be noted that it may overestimate the frequency of losses. For example, GBM is a commonly used choice in the industry when quickly hedging and pricing common vanilla options with high liquidity.

When markets are highly volatile such as during crises, or when a detailed assessment of product tail risk is required, GARCH models are better able to capture the impact of extreme events. Although they predict losses less frequently, they reveal the severity of losses once they occur, which is crucial for risk management. For example, when stress testing or calculating an insurance company's solvency capital, using GARCH models can more accurately reflect potential losses under extreme market scenarios.

For path-dependent products like automatically redeemable notes, since price performance on the observation date is crucial, the time-varying volatility of GARCH can cause the path's volatility characteristics to differ before and after the observation date, thus affecting the redemption probability. Therefore, GARCH is recommended for precise analysis. For example, in the Trigger Autocallable Notes studied in this paper, redemption depends on whether the price reaches a threshold on each observation date. GARCH can better simulate volatility changes around the observation date.

4. Conclusion

This paper takes an automatically redeemable note issued by HSBC and linked to the S&P 500 index as an example. It uses both Geometric Brownian Motion (GBM) and GARCH (1,1) models to perform Monte Carlo simulations of the underlying asset's price path, comparing the product valuation and risk indicators of the two models within a risk-neutral pricing framework. The conclusion is that there are significant differences in the fair value estimation and risk characteristics of the two models. The GARCH model has a higher average return, indicating that considering the time-varying nature of volatility improves the theoretical value of the product. The risk of GBM is mainly manifested in a higher frequency of losses and a larger expected loss, while the risk of GARCH is reflected in a lower frequency of losses but still considerable expected losses, showing the occasional extreme losses under time-varying volatility. GBM's constant volatility assumption ignores volatility clustering and mean reversion, while GARCH, by characterizing volatility dynamics, makes the path closer to the mean reversion characteristics of the actual market, providing a more refined risk-return profile. Therefore, the GBM model is more suitable for scenarios with relatively stable markets, no obvious volatility clustering, or simple product structures (such as European options). The GARCH model is more suitable for scenarios with high market volatility, volatility clustering, path-dependent products (such as automatically redeemable notes and cumulative options), or where precise assessment of tail risk is required.

Based on the empirical results and model comparisons of this study, several limitations need to be pointed out: First, the risk-neutral framework assumes zero drift and a normal distribution for GARCH residuals, failing to fully capture the skewed and fat-tailed characteristics of actual returns. Second, the sample period is limited to 2021–2024, and the product structure only involves auto-redeemed notes linked to the S&P 500; thus, the generalizability of the conclusions needs further verification. To address these shortcomings, future research is recommended to introduce EGARCH or asymmetric volatility models, adopt t-distribution or GED (Growth Equal Distribution)

DOI: 10.54254/2754-1169/2026.33934

assumptions for fat-tailed residuals, and extend the scope to multi-asset linked products or products with different market cycles. Looking ahead, deep learning time-series models could be combined to generate price paths, or GARCH tail loss estimation could be applied to capital allocation and stress testing, thereby further improving the accuracy and practical applicability of structured product pricing and risk management.

References

[1] Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81(3), 637–654.
[2] Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. Econometrica, 50(4), 987–1007.
[3] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307–327.
[4] Mozumder, S., Talukdar, B., Kabir, M. H., & Li, B. (2024). Non-linear volatility with normal inverse Gaussian innovations: Ad-hoc analytic option pricing. Review of Quantitative Finance and Accounting, 62, 97–133.
[5] Pang, T., & Zhao, Y. (2025). On GARCH and autoregressive stochastic volatility approaches for market calibration and option pricing. Risks, 13(2), 31.
[6] Sidhu, K. S., & Saxena, P. (2026). Beyond Black-Scholes: A computational framework for option pricing using Heston, GARCH, and jump diffusion models [Preprint]. arXiv.
[7] Hansen, P. R., & Tong, C. (2026). Option pricing with time-varying volatility risk aversion. The Review of Financial Studies, 39(3), 875–924.
[8] Wang, C., & Gerlach, R. (2023). A Bayesian realized threshold measurement GARCH framework for financial tail risk forecasting. Journal of Forecasting, 43(1), 40–57.
[9] Cui, Y., Li, L., & Zhang, G. (2024). Pricing and hedging autocallable products by Markov chain approximation. Review of Derivatives Research, 27, 259–303.
[10] Cheng, J. (2025). Enhancing S&P 500 index option pricing with GARCH model: A study on AI-induced market volatility. In Proceedings of the 2025 7th International Conference on Economic Management and Cultural Industry (ICEMCI 2025) (pp. 394–407). Atlantis Press.