Proceedings of ICMRED 2026 Symposium: The Future of Work: Strategy, Workforce Transformation, and Organizational Renewal
DOI: 10.54254/2754-1169/2026.LD31751

U.S. Short-Term Interest Rate Path Simulation and Derivatives Pricing: An ARIMA-GARCH Approach

Siyi Sun

Department of Mathematics, University College London, London, United Kingdom
zcahss4@ucl.ac.uk

Abstract. This study develops and evaluates a hybrid ARIMA–GARCH model with Student-t innovations to simulate the U.S. short-term interest rate path and price the interest rate caplet. As a key monetary policy tool and a fundamental benchmark in financial markets, the behaviour of short-term interest rates is characterised by volatility clustering and non-Gaussian innovations, features that constant-volatility lognormal models inadequately capture. Utilising the U.S. 3-Month Treasury Bill rate as a proxy for the short rate, this study estimates the model parameters and simulates 10,000 Monte Carlo paths over a three-month horizon. A caplet is then valued under a short-rate framework with payoff $N\tau max(r_T - K, 0)$, and pathwise discounting, and the resulting price is compared with a Black–76 benchmark. The empirical results indicate strong volatility persistence and heavy-tailed innovations, supporting the need for time-varying volatility and non-Gaussian errors. In pricing, the ARIMA–GARCH–t Monte Carlo approach yields a significantly higher caplet value than the Black–76 benchmark, underscoring the sensitivity of derivative values to volatility and distributional assumptions. Overall, the proposed framework provides a more realistic representation of short-rate risk and a more defensible basis for caplet valuation, with practical value for hedging and stress testing in interest-rate markets.

Keywords: Short-Term Interest Rates, ARIMA-GARCH Model, Monte Carlo Simulation, Caplet Pricing, Volatility Clustering

1. Introduction

1.1. Research background

Short-term interest rates are considered to be central to monetary policy transmission and also serve as a core reference point in the financial market. Practically, the U.S. 3-month Treasury bill yield is widely treated as a near risk-free short-rate benchmark due to its very low credit risk, short maturity and transparent prices. Therefore, it is often used as a key input in valuing and discounting a broad range of instruments, from household borrowing products such as mortgages to more complex interest-rate derivatives. However, in the post-crisis era, short-term interest rates have become harder to model, as shifts in monetary policy, inflation pressures, and geopolitical uncertainty have made their time-series behaviour less stable [1]. A prominent feature observed in daily interest rate data is "volatility clustering," meaning that large changes in interest rates tend to be followed by

DOI: 10.54254/2754-1169/2026.LD31751

large changes, and small changes by small ones [2]. In other words, it implies that volatility is time-varying instead of constant. Such behaviour directly affects financial institutions' risk management, as traditional pricing models, which assume constant volatility, often fail to capture these dynamics, creating challenges for accurate valuation and hedging. To manage this short-rate risk and reduce valuation and hedging errors, market participants often use interest rate derivatives like caplets, which are widely used as an insurance policy that pays off when the reference rate exceeds a pre-determined strike rate, to protect against unexpected increases in short-term rates. For example, commercial banks and hedge funds rely heavily on interest rate derivatives, such as Interest Rate Caplets, in order to hedge against the risk of rising borrowing costs. Thus, because of the vital role of caplets in the interest rate market, accurate pricing is commercially critical, helping financial institutions to minimise the risk. Mispricing may trigger arbitrage losses or hedging failures. Therefore, creating an accurate pricing model to reflect the nature of interest rates is not only an academic undertaking but also a practical requirement in risk management in the contemporary banking industry.

1.2. Literature review

The literature review proceeds in three parts: (1) examines the statistical properties of market rates; (2) discusses distributional assumptions; and (3) reviews derivative valuation methods. Recent advances in financial econometrics have increasingly challenged the overly idealistic assumptions (e.g., constant volatility and approximately normal errors) of traditional interest rate models, particularly since the post-2020 period, when the high-volatility environment began. Recent literature has shown a shift towards hybrid models that can capture stylised facts such as volatility clustering and fat tails.

Firstly, in the context of the statistical characteristics of the market rate, the article by Enow focused on the dynamics of volatility in the financial market from 2018-2023. He discovered that volatility clustering is still an eminent phenomenon in which the large fluctuations are continuously succeeded by large fluctuations [3]. This observation justifies the application of GARCH models in comparison to constant variance models since the latter model is inadequate in explaining time-varying risk that currently prevails in the current economic landscape. His discovery is the direct motivator of the selection of GARCH modelling for the interest rate volatility. Besides, Dean and Francis suggested an ARIMA-GARCH-S model of short-term stock prediction and showed that the hybrid model beats the linear models in the short-term forecasting of stock, both in capturing the mean trend and volatility that varies over time [4].

In addition to that, with regard to the distributional assumptions of simulating the path, it has been recently stated that the normality assumption in the Black-76 model is not acceptable. Wang compared the option pricing of Monte Carlo simulations and discovered that the traditional price of options, including the Black-Scholes model, tends to have a large pricing error since they are unable to model the "volatility of volatility" and the jump-diffusion processes in the real market [5]. Studies by Nugroho et al., therefore, suggested that the application of t-distribution in GARCH modelling would best fit the financial returns because the t-distribution will help locate the asymmetry and heavy tails that have not been adequately captured by normal models [6]. Besides, GARCH-type models with a Student t-distribution were additionally found to better fit bond markets in the volatility clustering and leptokurtosis-related models, compared to models assuming Gaussian errors [7]. This finding rejects the standard normal hypothesis and presents empirical evidence of the use of the t -distribution in the GARCH model applied in the current study. Moreover, regarding the derivative pricing, Zheng analysed the shortcomings of the analytical formulas in the path-

DOI: 10.54254/2754-1169/2026.LD31751

dependent options and found that Monte Carlo can be used to price more complicated derivatives, including caps, due to the capability of discrete monitoring and dynamic changes in volatility [8]. In addition, recent research has highlighted the limitations of the classic diffusion models in the discrete movements of interest rates. As an example, da Silva and Baczynski have shown that the use of discrete jumps and scheduled central bank interventions in pricing models is a great enhancement of the accuracy of valuing interest rate derivatives over continuous models [9].

1.3. Research gap

In spite of these developments, there are still a number of key gaps that exist regarding the practical use of these models. Although the current literature has dealt extensively with interest rate modelling and derivative pricing, there are some gaps in practice. The continuous-time models, including Vasicek models and the CIR model, have been proposed as a research subject and are mostly reviewed by scholars, or the use of the Black-Scholes model in pricing caps of interest rate. The hybrid ARIMA-GARCH model with Student-t distributions to value interest rate derivatives through Monte Carlo simulation has hardly been used in any research. In particular, the application of ARIMA to the mean dynamics, GARCH to the volatility clustering and t-distribution to the heavy tails has never been systematically applied to Caplet pricing.

On the contrary, the paper finds that despite the popularity of the GARCH models in equity markets, their ability to rectify pricing biases in the Interest Rate Caps market is not fully researched. In an attempt to overcome the weaknesses of the normal assumption, the aim of the research is to come up with a more acceptable pricing alternative, that is, one that considers the extreme risks associated with the market. The need to fill this gap is especially urgent due to the recent revival of interest rate instability in post-pandemic markets.

1.4. Research framework

This paper seeks to formulate an effective pricing model of Interest Rate Caplets that takes into consideration both the non-normalised volatility as a function of time, and non-normal distributions. Firstly, statistical analysis of the U. S. 3-Month Treasury Bill Secondary Market Rate is performed in the paper. Based on the observed high autocorrelation and volatility clustering, an ARIMA-GARCH model is developed. The ARIMA component is used to reflect the average trend, whereas the GARCH component depicts the volatility, which changes over time. Secondly, this paper uses a Student-t distribution as an addition to the residual of the GARCH model, which will correct the "thin tail" bias of the conventional Gaussian established models. This would be a major part of capturing the fat tail properties of interest rate change, so that the likelihood of extreme events in the market can not be underestimated. The Monte Carlo method offers the capability of modelling path-dependent characteristics, unlike closed-form solutions that rely on restrictive conditions. Third, it is possible to conduct a Monte Carlo simulation with an estimate of 10, 000 possible future interest rate values based on the estimated parameters of the ARIMA-GARCH-t model so that the payoff of the Interest Rate Cap at maturity can be computed. After all, the fair caplet value is obtained by discounting payoffs path by path under the simulated short rate and then averaging across simulations. The obtained results are then contrasted with the prices derived from the standard Black-76 model by calculating the difference between the benchmark Black-76 and the ARIMA-GARCH model price. This comparative analysis not only validates the model but also quantifies the pricing biases inherent in traditional approaches.

DOI: 10.54254/2754-1169/2026.LD31751

2. Methodology: data, model specification, and simulation framework

2.1. Data and descriptive statistics

In this study, secondary data retrieved from the Federal Reserve Economic Data (FRED) database will be used to analyse pricing of interest rate derivatives, specifically the daily 3-Month Treasury Bill Secondary Market Rate, which serves as a representation of the short-term risk-free rate (rt). This study implements all computations in Python, using the statsmodels library for stationarity tests (ADF) and autocorrelation analysis (ACF/PACF), and the arch library for volatility modelling. Before model estimation, this study conducted preliminary data analysis to ensure stationarity and identify appropriate lag structures.

2.2. Model specification: ARIMA–GARCH–t framework

To capture the characteristics of financial time series, the study constructs a hybrid ARIMA-GARCH model. Given the volatility clustering, this study adopts a GARCH framework to model time-varying volatility. Firstly, the ARIMA component is used to model the linear dependence and mean trend of the interest rate series (E[rt|It-1]), where the optimal lag order (p, d, q) is determined based on the Akaike Information Criterion (AIC) to ensure the model parsimony. The GARCH(1,1) component models the time-varying conditional volatility (Var[rt|It-1]). Specifically, the variance equation determines current volatility based on past market shocks (ARCH term) and previous volatility levels (GARCH term). While the GARCH terms capture the magnitude of market volatility, the t distribution models the probability of extreme occurrences within that volatile environment. By estimating the degrees of freedom parameter ν, the model quantifies the extent of heavy tails (leptokurtosis) in interest-rate changes, reducing the underestimation of extreme moves implied by a Gaussian assumption.

2.3. Monte Carlo simulation and pricing algorithm

Following this, the simulated interest rate directions are generated using a Monte Carlo model with an estimated parameter. 10,000 simulated interest rate paths were created using a 3-month horizon that comprises 63 trading days. This Monte Carlo method has two major advantages compared to closed-form solutions: unlike other standard models, including the Black-Scholes model, this simulation is recursive. The volatility is modelled dynamically in every time period, and it is given the prior day's simulated shock. Therefore, this method would have greater flexibility than the Black-76 model by capturing the path-dependent nature of volatility accumulation.

Finally, the caplet value is determined by considering 10, 000 simulated short-rate paths during a 3-month period and the value of the payoff at maturity. The present value is obtained by discounting payoffs along each simulated path under a short-rate discount factor, which is a discrete-time approximation to continuous compounding, and then averaging the discounted payoffs across all paths. Then, the model results are strictly compared with the industry standard Black-76 model to ensure that the model works well. Such competitive analysis reveals the exact pricing difference caused by incorporating stochastic volatility and heavy-tailed distributions. After that, to ensure robustness, this study conducts sensitivity analyses by varying the number of simulations and the initial volatility level.

DOI: 10.54254/2754-1169/2026.LD31751

3. Empirical results: model estimation and pricing performance

This chapter provides the empirical knowledge of the quantitative structure of the methodology. The system analysis is done in a systematic way. To begin with, the statistical characteristics of the simulated data of the 3-month Treasury Bill, which is used to prove the use of the ARIMA-GARCH method, are discussed. Second, the chapter provides the results of the estimation of the hybrid model, emphasising the statistically significant parameters. Third, the interest rate caplet outcome of pricing is compared to the Black-76 benchmark. Lastly, the paper points out the ways in which conventional pricing models fail.

Besides, this chapter offers empirical information of the DTB3 time series as a test of the suggested ARIMA–GARCH–t model. The analysis follows the following way. The U.S. 3-Month Treasury Bill rate observed is analysed first to stimulate the ARIMA–GARCH specification. Second, the results of the estimation are described and discussed, taking into consideration statistical and economic significance. Lastly, the caplet price of the ARIMA-GARCH-t Monte Carlo model is compared against the Black-76 benchmark.

3.1. Data diagnostics and model justification

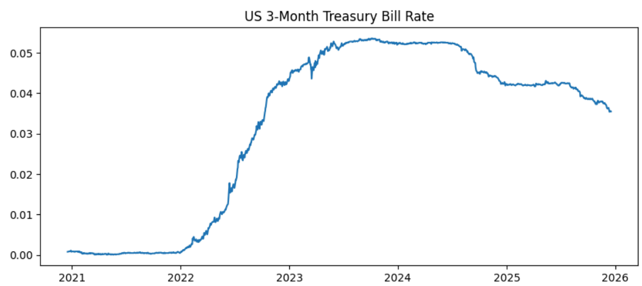

Figure 1. U.S. 3-month treasury bill rate

(Source: Federal Reserve Economic Data (FRED), series DTB3)

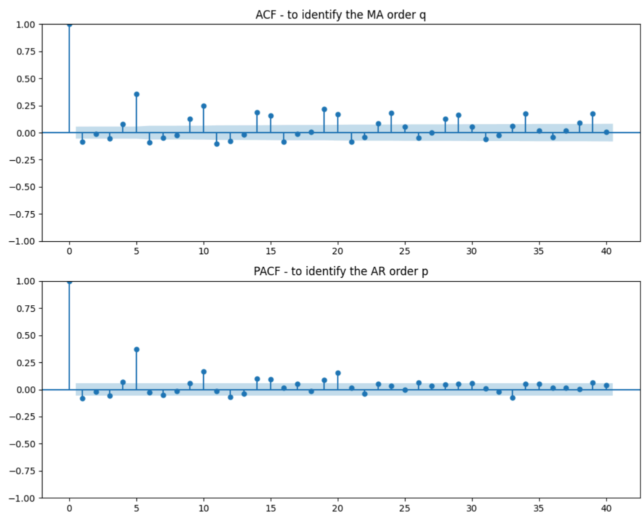

Figure 2. ACF and PACF of the first-differenced series

(Source: The study's calculations using FRED (DTB3))

DOI: 10.54254/2754-1169/2026.LD31751

Figure 1 plots the 3-month Treasury bill rate in levels. The series is highly persistent and moves through clear phases rather than fluctuating around a stable mean. This visual pattern suggests that the level series may be non-stationary. To test this formally, Augmented Dickey–Fuller (ADF) tests are applied. This study performs the ADF test with a constant and trend term, using the AIC to select the optimal lag length.

For the level series, the ADF statistic is −1.5832 with p = 0.492, so the unit-root null is not rejected, which means it contains a unit root and is not stationary. Therefore, the first difference is used. After taking the first difference, the ADF statistic becomes −4.2036 with a p-value of 0.0006496, which strongly rejects the unit-root null, indicating the stationarity of the series. Therefore, the short rate is treated as I(1), meaning it becomes stationary after one differencing step. This directly motivates using a d = 1.

After differencing, the next question is whether changes in the rate are close to white noise or whether they still contain short-run structure. Figure 2 shows the ACF and PACF of the first-differenced series. Instead of dying out immediately, both functions display noticeable early-lag spikes, which indicates that the differenced series still contains short-run dependence. Therefore, the mean equation should include AR and MA terms (an ARMA structure in differences), rather than assuming the differenced series is random noise. This motivates estimating an ARIMA(p, 1, q) model for the level rate.

3.2. Estimation of the ARIMA–GARCH–t model

Given the stationarity results, candidate ARIMA models with d = 1 are compared using the Akaike Information Criterion (AIC). The selected model is ARIMA(5,1,5) with no deterministic trend (trend = "n"), yielding an AIC of −16513.634 (reported in Table 1). This selection is consistent with the dependence pattern in Figure 2, which suggests that more than one or two lags may be needed to capture the short-run dynamics in rate changes.

Table 1. Estimated ARIMA–GARCH model parameters

Panel A: ARIMA(5, 1, 5) mean equation (trend = none)			
Parameter	Estimate	Std. Error	p-value
AR(1)	-0.0606	0.009	0.000
AR(2)	-0.0833	0.008	0.000
AR(3)	-0.0621	0.008	0.000
AR(4)	0.1995	0.007	0.000
AR(5)	0.7784	0.009	0.000
MA(1)	-0.0515	0.010	0.000
MA(2)	0.1303	0.014	0.000
MA(3)	0.0469	0.010	0.000
MA(4)	-0.1289	0.016	0.000
MA(5)	-0.4938	0.014	0.000
AIC		-16513.634	
nobs		1250	
Panel B: GARCH(1,1)-t variance equation			
	Estimate	Std. Error	p-value
ω	2.0679e-05	3.863e-07	0.000
α1	0.1980	2.352e-02	3.886e-17
β1	0.7722	2.718e-02	1.431e-05

DOI: 10.54254/2754-1169/2026.LD31751

Table 1. (continued)

α1+β1	0.9702	-	-
ν (t d.f.)	4.47	-	-

Table 2. Ljung–box test on ARIMA residuals

Lag	Ljung–Box statistic	p-value
10	6.071	0.809
20	31.713	0.046
30	47.790	0.021

Table 3. ARCH-LM test on ARIMA residuals

Test	Statistic	p-value
ARCH-LM (LM)	142.697	< 0.001
ARCH-LM(F)	15.982	< 0.001

After fitting an ARIMA model, residual checks are performed to understand what remains unexplained. A key aim of the ARIMA stage is to remove predictable mean dynamics, so that the residuals behave like unpredictable shocks with minimal linear dependence. The Ljung–Box results (Table 2) suggest that the ARIMA specification removes much of the linear dependence in the mean, especially at shorter lags. However, the p-values at lags 20 and 30 are below 0.05, indicating that some residual autocorrelation may remain at longer horizons. But this does not invalidate the mean model; it motivates cautious interpretation and supports focusing on conditional heteroscedasticity in the next step. To test this, an ARCH-LM test is applied to the ARIMA residuals (Table 3). The test yields an extremely small p-value, providing strong evidence against the assumption of constant variance. In economic terms, this means the magnitude of shocks changes over time: quiet periods tend to be followed by quiet periods, and volatile periods tend to be followed by volatile periods. This is precisely the volatility clustering pattern that motivates a GARCH variance model.

To capture this behaviour, a GARCH(1,1) model is fitted to the ARIMA residuals, with Student-t errors. The GARCH(1,1) model is used because it is the standard, parsimonious specification that can effectively model volatility clustering. Student-t errors are also adopted rather than normal errors, because interest-rate changes can occasionally contain large, rare moves. The Student-t distribution has heavier tails, assigning more probability to extreme outcomes and typically fitting financial data better.

The estimated GARCH parameters (reported in Table 1) are statistically significant. In particular, the sum of the ARCH and GARCH coefficients satisfies

$$\alpha 1 + \beta 1 \approx 0.1980 + 0.7722 \approx 0.97 \tag{1}$$

indicating a high degree of persistence in the conditional variance, such that the impact of volatility shocks diminishes gradually over time. In addition, the estimated degrees of freedom are $\nu \approx 4.47$, indicating heavy tails relative to a normal distribution.

Overall, the estimation results support the hybrid structure: ARIMA captures the short-run mean dynamics in rate changes, while GARCH captures clustered, persistent volatility, and the Student-t errors further allow for occasional large shocks. This empirically validates the core modelling premise of this study: that a hybrid ARIMA-GARCH framework with Student-t innovations is

DOI: 10.54254/2754-1169/2026.LD31751

necessary to adequately capture the persistence, volatility clustering, and tail risks inherent in the short-term interest rate series, a combination often overlooked in standard pricing models like Black-76.

3.3. Caplet pricing and comparative analysis

This section compares caplet prices obtained from the ARIMA–GARCH–t Monte Carlo framework with a Black–76 benchmark under constant volatility. Figure 3 shows a subset of short-rate paths simulated via Monte Carlo under the estimated ARIMA–GARCH–t model. Most paths fluctuate around the current level, but occasional sharp movements also appear, which is consistent with time-varying volatility and heavy-tailed shocks. These features are important for caplet valuation because the payoff depends on whether the rate ends above the strike at maturity. Although large upward moves are not frequent, they would have a noticeable impact on the expected payoff.

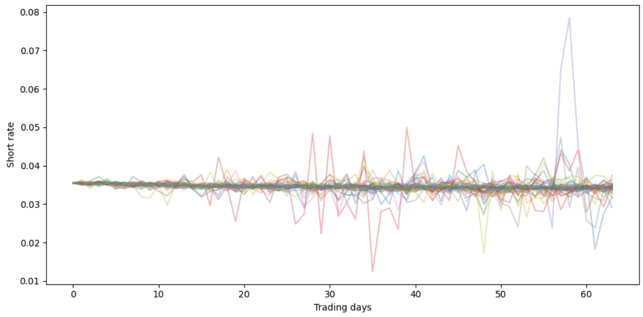

Figure 3. Simulated short-rate paths under ARIMA–GARCH–t

(Source: This study's calculations using FRED (DTB3))

Table 4. Caplet pricing comparison

Model	Caplet price
ARIMA-GARCH-t	82.71
Black-76	16.55

Using the simulated paths, the caplet payoff at maturity T is specified as

$$Payoff = N \times \tau \times max(r_T - K, 0) \tag{2}$$

where N denotes the notional principal and τ is the year fraction of the underlying interest period. The term rT represents the reference short-term interest rate at the settlement date T generated from the simulated interest-rate path, and K is the contractual strike rate. The function $max(r_T - K, 0)$ reflects the option feature of a caplet: the payoff is positive only when the realised rate exceeds the strike, and is zero otherwise.

To obtain the present value, payoffs are discounted along each simulated path. A discrete-time approximation is used to represent continuous-compounding discounting under the short-rate framework:

$$\mathrm{D}\left(0, T\right) \approx \exp\left\{-\sum_{t=1}^{T} r_t \times \Delta t\right\} \tag{3}$$

DOI: 10.54254/2754-1169/2026.LD31751

where rt is the simulated short rate at time step t along the path, and Δt is the time increment in years (for daily data, typically $\Delta t = \frac{1}{252}$). The summation $\sum_{t=1}^{T} r_t \times \Delta t$ serves as a Riemann-sum approximation to the integral $\int_0^T r_s ds$. Therefore, the discount factor for each path is computed as $\exp\left\{-\sum_{t=1}^{T} r_t \times \Delta t\right\}$, which reflects the cumulative effect of the simulated short rates over the time period $[0,\ T]$.

Under the ARIMA–GARCH–t model, the Monte Carlo price is 82.71 (Table 4). For comparison, Black–76 is applied with a constant volatility estimate. Let

$$\sigma_{abs} = \sqrt{252} sd\left(\Delta r_t\right) \tag{4}$$

denote the annualised volatility of daily changes in the short rate. This is converted to the lognormal-style volatility used in Black–76 by normalising with the current rate level F0, i.e.

$$\sigma_{Black} = \frac{sd(\Delta r_t)\sqrt{252}}{F_0} \tag{5}$$

yielding a Black–76 price of 16.55 (Table 4).

4. Discussion

4.1. Principal empirical findings and pricing discrepancy

This paper hypothesises and tests empirically a hybrid ARIMA-GARCH-t policy of caplet pricing, where it is compared to the conventional Black-76 model. Its main result is a huge and economically substantial pricing difference: the ARIMA-GARCH-t Monte Carlo 82.71 is compared to the Black-76 16.55. This is a difference that implies that under time-varied risk and non-Gaussian shocks, a constant-volatility lognormal benchmark overvalues options. Specifically, adding conditional heteroscedasticity and heavy-tailed innovations enhances the chances of having large positive changes in rates during the horizon and in this way, the caplet's expected payoff and its present value after discounting increase.

4.2. Sources of the pricing discrepancy: a tripartite advantage of the ARIMA–GARCH–t framework

The ARIMA–GARCH–t framework offers three clear advantages over constant-volatility lognormal benchmarks in caplet pricing.

First, it captures time-varying risk through GARCH dynamics. With $\alpha 1 + \beta 1 \approx 0.97$, volatility is highly persistent, so the model can reflect the empirical feature that uncertainty remains persistently high after shocks. This makes the distribution of future rates wider in stressed periods, whereas a constant σ benchmark cannot adjust and therefore tends to understate tail risk, which directly contributes to a higher probability of the underlying rate exceeding the strike price (K), thereby increasing the expected payoff in the Monte Carlo simulation.

Second, it accommodates non-Gaussian shocks via Student-t innovations. The estimated degrees of freedom $\nu \approx 4.47$ indicate pronounced heavy tails, so extreme short-rate changes are more likely than under Gaussian errors. This matters for caplets because the payoff $N\tau max(r_T - K,\ 0)$ is driven by the right tail of rT. If the model understates the probability of

DOI: 10.54254/2754-1169/2026.LD31751

large upward rate moves, it will systematically underestimate the frequency and magnitude of positive payoffs. By assigning greater probability to extreme outcomes, the Student-t specification produces a more credible payoff distribution and, therefore, a more reliable caplet valuation than a Gaussian benchmark.

Third, the Monte Carlo implementation provides a coherent valuation procedure for this model. Payoffs $N\tau max(r_T - K, 0)$ can be evaluated directly on each simulated path, and discounting can be applied consistently within a short-rate framework using a pathwise discount factor (a discrete approximation to continuous compounding). This simulation-based approach does not require a closed-form pricing formula. It therefore remains practical when the model includes stochastic volatility and heavy-tailed errors, and it generalises naturally to more complex payoffs where analytic benchmarks are unavailable.

4.3. Alignment with and extension of existing literature

The findings are in agreement and consistent with various threads of previous research. First, the presence of volatility clustering, as well as high variance persistence, is consistent with the findings provided by Enow which demonstrate that despite the fact that GARCH-type structures are appropriate for interest-rate dynamics rather than being restricted to equity returns [3]. This finding is also supported by Pastpipatkul and Ko, who showed that time-varying volatility models remain the superior method to study the returns of financial assets in a complex market space [10]. Second, the estimated heavy tails and their improvement compared to the Gaussian standards are in line with those of Nugroho et al., who contend that t-innovation can more favourably model the extreme movements and reduce systematic underestimation of tail risk [6]. Third, the practical feasibility of the pricing exercise supports Zheng, through the flexibility of Monte Carlo methods when dealing with models with no closed-form pricing and with state variables (such as conditional variance) that are dynamic [8].

4.4. Practical and policy implications for risk management

In a risk-management sense, the scale of pricing difference implies that relying on constant-volatility benchmarks may leave financial institutions exposed to systematic under-hedging. The banks and hedge funds deploying caplets to address funding cost risk must have valuations that are strong even at times of volatility. As an example, a trading desk, working with this framework of the study, would be able to more accurately price and execute a caplet position more efficiently during a regime of declared monetary policy uncertainty, as the model would automatically reflect the higher volatility and tail risk in the price. A time-varying volatility model in this case gives a more responsive contribution to prices and hedges. For regulators, models that account for time-varying volatility and heavy tails can enhance the evaluation of interest-rate risk, particularly when assessing how portfolios perform under adverse scenarios. Stress tests based on simulated paths from this study's model, which include clustered high-volatility regimes and extreme jumps, would provide a more severe yet realistic assessment of capital adequacy than those based on constant-volatility assumptions. More broadly, improved pricing accuracy can enhance market efficiency by reducing mispricing-driven arbitrage and by aligning derivative prices more closely with the underlying risk environment.

DOI: 10.54254/2754-1169/2026.LD31751

5. Conclusion

5.1. Summary of core findings

This study explored whether a hybrid time-series framework can value caplets more effectively than standard constant-volatility benchmarks. Using the U.S. 3-Month Treasury Bill rate as a measure of the short rate, an ARIMA–GARCH model with Student-t innovations was fitted and then used in a Monte Carlo simulation to generate forward short-rate paths. These simulated paths were used to price a caplet with payoff $N\tau max(r_T - K, 0)$, with payoffs discounted using the simulated short-rate paths.

By using these methods, three main findings emerge. First, the short-rate series exhibits strong persistence and clear changes in volatility over time, which supports the use of a conditional heteroscedasticity model rather than assuming constant variance. Second, the estimated innovation distribution is heavy-tailed, meaning that extreme rate movements are more likely than under Gaussian errors. Third, Monte Carlo valuation is well suited here: it updates the variance over time along each simulated path, evaluates the nonlinear caplet payoff directly, and does not rely on a single analytic pricing formula. Consequently, the proposed framework generates a caplet price approximately five times higher than the Black-76 benchmark (82.71 vs. 16.55). This substantial discrepancy highlights that ignoring volatility clustering and tail risk results in a severe undervaluation of interest rate options, underscoring the material economic impact of model specification.

5.2. Theoretical and practical contributions

This study is of practical interest in rate derivative pricing and risk management. It gives a closer image of the short-rate risk by modelling time-varying volatility and heavy-tailed shocks in comparison with constant-volatility Gaussian benchmarks. This justifies more defensible caplet valuation, and makes institutions better run funding-cost exposure. The Monte Carlo framework is also helpful with stress testing, as unfavourable situations can be analysed using simulated short-rate paths. In more general terms, the study supports the motivation of the topic: as short-term rates tend to exhibit volatility clustering and can generally rapidly alter when policy or market structure changes, the valuation of derivatives ought to be based on these empirically observed short-rate patterns rather than simple-minded assumptions.

5.3. Limitations and future research directions

However, several limitations remain. First, parameter estimates may be sensitive to the sample window. Short-rate dynamics can change with the macro and policy environment, so coefficients fitted in one period may not be stable in another, which limits how far the results can be generalised. Future research could employ rolling-window or recursive estimation to test the stability of these parameters across different monetary policy regimes. Second, ARIMA assumes the same linear pattern holds over time. However, policy or expectation shocks can cause abrupt changes in the overall rate level and its persistence, which a fixed-order ARIMA does not explicitly capture. Incorporating Markov-switching or time-varying parameter (TVP) extensions into the mean equation could be a fruitful avenue to capture such structural breaks. Third, the Black–76 benchmark depends on volatility calibration. Because Black–76 requires a single constant σ, different choices of how σ is estimated can materially change the benchmark price, affecting the interpretation of pricing differences.

Building on this work, promising future directions include: (1) investigating regime-switching GARCH or more general heavy-tailed distributions (e.g., skewed-t) to capture asymmetric market phases better; (2) extending the framework to a multi-curve or multi-factor setting to ensure consistency across the entire term structure; and (3) conducting out-of-sample hedging performance tests to compare the economic utility of the proposed model against traditional benchmarks.

References

[1] International Monetary Fund. (2013). Unconventional monetary policies: Recent experience and prospects (IMF Policy Paper). International Monetary Fund. https: //www.imf.org/external/np/pp/eng/2013/041813a.pdf
[2] Koedijk, C. G., Nissen, F. G. J. A., Schotman, P. C., & Wolff, C. C. P. (1997). The dynamics of short-term interest rate volatility reconsidered. European Finance Review, 1(1), 105–130.
[3] Enow, S. T. (2023). Exploring volatility clustering in financial markets and its implication. Journal of Economic and Social Development.
[4] Dean & Francis. (2024). Application of ARIMA-GARCH-S model to short-term stock forecasting. Dean & Francis Press. https: //www.deanfrancispress.com/index.php/fe/article/download/271/FE000502.pdf
[5] Wang, Y. (2024). Research on pricing of SSE STAR 50ETF options based Monte Carlo simulation. Advances in Economics, Business and Management Research. Atlantis Press. https: //www.atlantis-press.com/article/126016392.pdf
[6] Nugroho, D. B., Setiawan, A., & Morimoto, T. (2025). Modelling and forecasting financial volatility with realised GARCH model: A comparative study of skew-t distributions using GRG and MCMC methods. Econometrics, 13(3), 33.
[7] Hadad, E., Fridman, A. M., & Yosef, R. (2025). Estimating corporate bond market volatility using asymmetric GARCH models. Risks, 13(11), 224.
[8] Zheng, S. (2022). Pricing barrier options with simulations. SimTrade. https: //www.simtrade.fr/blog_simtrade/tag/monte-carlo-simulations/
[9] da Silva, A. J., & Baczynski, J. (2024). Discretely distributed scheduled jumps and interest rate derivatives: Pricing in the context of central bank actions. Economies, 12(3), 73.
[10] Pastpipatkul, P., & Ko, H. (2025). Volatility analysis of returns of financial assets using a Bayesian time-varying realised GARCH-Itô model. Econometrics, 13(3), 34.