Forecasting returns volatility : A GARCH-type approach on NVDA (2018-2023)

Gabriel AMMOUR* and Noa LE ROUX*

*2^{nd} year students in Applied Econometrics, Nantes Université

March 21, 2025

Abstract

This study examines the volatility forecasting of Nvidia (NVDA) stock returns from 2018 to 2023 using GARCH-type models under normal and Student distributions. By estimating and comparing the performance of GARCH, GJR-GARCH, IGARCH, and RiskMetrics models, we assess their predictive capabilities using standard forecasting accuracy metrics. Our results indicate that while the GJR-GARCH model slightly outperforms others, the Diebold-Mariano and Model Confidence Set tests suggest no significant differences in predictive accuracy across models. These findings emphasize the robustness of various GARCH-type models in financial forecasting.

Keywords: Volatility forecasting, GARCH models, Nvidia, Financial econometrics, Stock market

1 Introduction

On 5 November 2024, the American company Nvidia (NVDA) overtook Apple and became the world's largest market capitalisation, with an estimated value of USD 3.4 trillion. The chip company is one of the big winners in the boom and commercial development of LLM artificial intelligence models such as ChatGPT. Just like cryptocurrencies, machine learning models require significant graphical computing power. According to the survey agency TrendForce, it would take more than 30,000 graphics cards to train the OpenAI model. Given that the graphics cards in question cost around USD 10,000 (Nvidia A100), we're talking about revenue of around 300 million for Nvidia just on this artificial intelligence model. Moreover, Nvidia is the world leader in mass-market graphics cards. With a market share of 90% in 2024, the Californian group enjoys a near-monopoly position, enabling it to charge higher prices than its direct competitors AMD and Intel.

Given the increasing importance of high-frequency trading and risk management in financial markets, accurately modeling and forecasting stock return volatility has become a critical challenge. Nvidia's dominant market position and exposure to technological advancements make it an interesting case for volatility analysis. The primary objective of this study is to determine which GARCH-type model provides the most accurate volatility forecasts for NVDA returns. To address this question, we will first describe the dataset and examine statistical properties of Nvidia's returns. Then, we will estimate various GARCH-family models under normal and Student distributions. Finally, we will compare forecasting accuracy through statistical tests, including the Model Confidence Set and Diebold-Mariano tests.

2 Data

2.1 Asset description

Founded in 1993 by current CEO Jensean Huang, Nvidia is a company that belongs to the large semiconductor industry. One must keep in mind that Nvidia is a software company which designs and supplies chips but does not make them. Similarly to many of its Californian competitors, Nvidia uses Asian suppliers to manufacture its products.

Nvidia went public on January 22, 1999 and the 3 largest shareholders of Nvidia in early 2024 were The *Vanguard Group* (8.280%), *BlackRock* (5.623%) and *Fidelity Investments* (5.161%). As mentionned earlier, Nvidia has one of the world's largest market capitalisation. Nvidia's asset can be traded on the NASDAQ, an American stock exchange based in New York City. It is the most active stock trading venue in the U.S. by volume, and ranked second on the list of stock exchanges by market capitalization of shares traded, behind the New York Stock Exchange. Additionnaly, the company is listed on the SP500 stock market index which is composed of 500 of the largest companies listed on stock exchanges in the United States.

In recent years, Nvidia has seen its sales and net profit grow significantly (Figure 1). In 2024, the company achieved record sales of US$60.92 billion and net income of US$29.76 billion. The company has more than 30,000 employees, which is 3 times more than 10 years ago in 2015 (Figure 1). This dynamic is particularly noticeable in the light of the massive downsizing that the tech world has experienced in the post-covid period.

Figure 1: Nvida revenue, net income and number of employees (2014-2023)

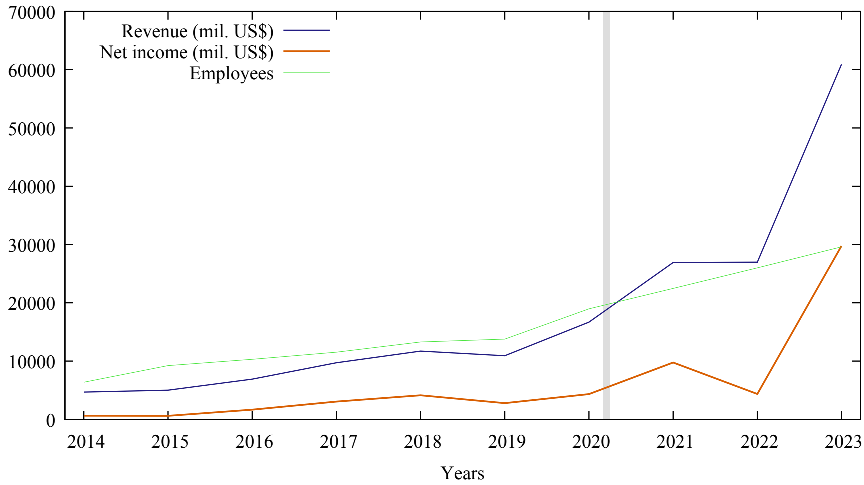

2.2 Outliers detection

Figure 2 shows the evolution of raw and corrected returns. Outliers can be identified by the green peaks. Among the detected points, three seem to be particularly important. The most important appears on 25 May 2023 and follows the announcement of the company's results. Following the boom in AI that year, the company made record profits and saw its share price rise significantly. The second outlier, this time negative, appeared on 16 March 2020. This fall was probably due to the uncertainty surrounding the lockdown and covid 19. On that day, the French President announced that France was going to war against the Covid19 epidemic. Finally, the third most important outlier occurred on 16 November 2018. After missing revenue expectations, the share price fell by 19%.

Figure 2: NVDA raw and clean returns over time (2018-2023)

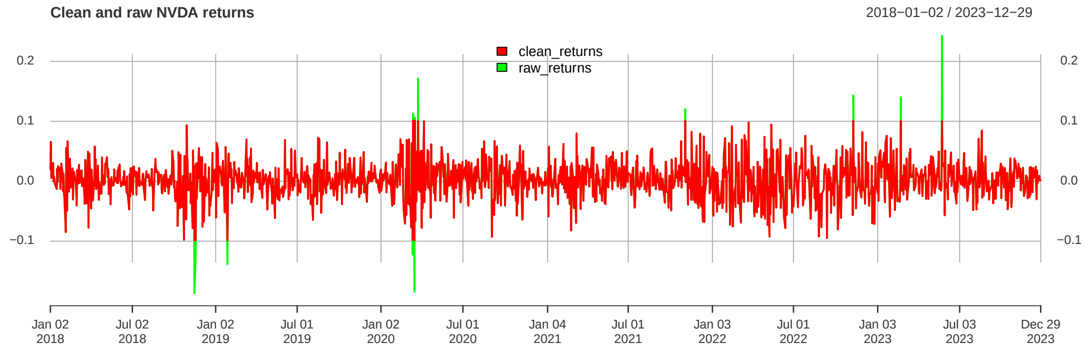

2.3 *Returns analysis*

Figure 3 shows the evolution of the NVDA share over time. Atypical values are identified by red dots. We can see that the share price rose sharply between 2020 and the end of 2023, from \$5 per share to \$50 per share, an increase of 1,000%. We can also see that the share price fell sharply throughout 2022, before rising again in 2023.

Figure 3: NVDA raw value over time (2018-2023)

Figure 4 shows the evolution of returns and returns squared for the nvidia share over the period studied. We notice that there is significant volatility, which is to be expected considering the daily data available. The squared returns provide us a proxy for the volatility of the share's returns.

Figure 4: NVDA clean returns and squared returns over time (2018-2023)

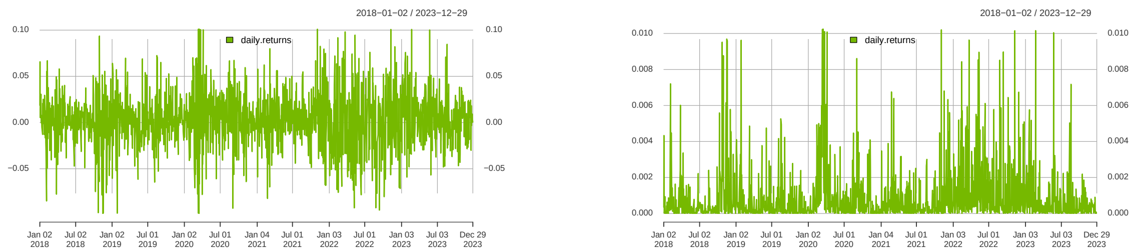

(a) Returns

(b) Squared returns

2.3.1 *Correlogram*

One of the properties of financial series is the autocorrelation of the squared returns, rt^2, whereas the profitability, rt, shows little or no autocorrelation. We are therefore interested in the autocorrelation and partial autocorrelation functions of the returns and the squared returns.

Figure 5 shows the ACF and PACF of the returns and the squared returns. Regarding the returns, we observe that the ACF shows only a very significant 1 lag and that the PACF has no definite lag structure. However, in terms of squared returns, we detect a decreasing structure in the number of lags, which indicates a certain persistence in volatility.

Figure 5: Returns ACF and PACF

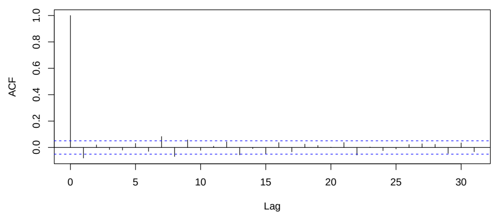

(a) Returns ACF

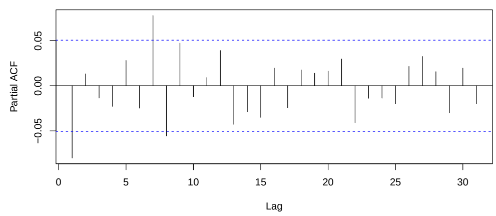

(b) Returns PACF

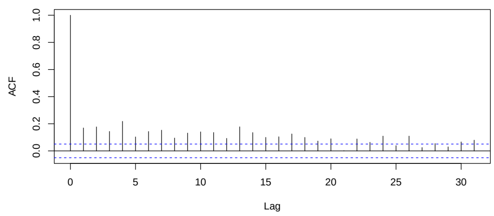

(c) Squared returns ACF

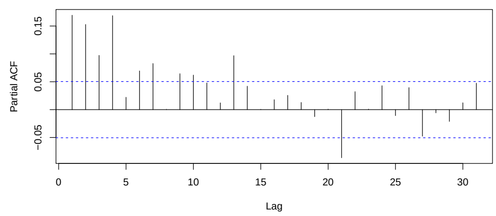

(d) Squared returns PACF

2.3.2 *Descriptive statistics*

Table I shows us the descriptive statistics for the NVIDIA share's returns, along with tests on the distribution of returns. The table shows that NVIDIA is a relatively profitable asset since the average and the median are both higher than the CAC40 reference asset. Over the period studied, the average daily return on the asset was 0.2 % compared with 0.03 % for the CAC40. In terms of volatility, measured by the standard deviation, the US asset is more volatile (3.06% > 1.18%) than the French index. This seems logical, since the CAC index comprises 40 different companies, and risk and volatility are mechanically lower. This increased volatility is confirmed by the maximum and minimum returns achieved by the NVIDIA asset, both of which are higher than those of the CAC40. NVDIA's extremes are around -10% and +10%, whereas the CAC40 is around -4.5% and +4.5%.

Concerning the distribution of returns, we note that the skewness is slightly below 0 (-0.019) and that the Kurtosis is above 0 (0.95). We can therefore expect the series not to follow a normal distribution. To confirm these intuitions, we can look at the statistical tests at the end of Table I. The Jarque-Bera test compares the skewness and kurtosis of the series studied with that of a Normal distribution. As the p-value of the test is less than 0.05, we reject the null hypothesis that the series follows a Normal distribution. Next, we used the Ljung-Box test to identify the presence or absence of autocorrelation. As the p-value of the test is less than 0.05, we reject the null hypothesis, which indicates that there is no autocorrelation between the returns. Finally, we completed our analysis by applying an LM-ARCH test for 5 and 10 lags. This test will allow us to identify potential heteroscedasticity in the residuals of the returns. For the two lags, the pvalue is less than 0.05, so we reject the null hypothesis which indicates that the residuals are homoscedastic. Following the three tests, we can conclude that our profitability series does not follow a Normal distribution and that the residuals are heteroskedastic and auto-correlated.

Table I: Descriptive Statistics and Tests for Daily Returns (Clean Data)

Descriptive Statistics	
Minimum	-9.84
Maximum	10.12
Mean	0.20
Median	0.26
Variance	9.37
Standard Deviation	3.06
Skewness	-0.019
v1	-0.31
Excess Kurtosis	0.95
v2	7.53
Statistical Tests	
Jarque-Bera (JB) Test	$X^2 = 57.495$ p-value $= 3.27 \times 10^{-13}$
Ljung-BoxQ(10)	$X^2 = 37.149$ p-value $= 5.33 \times 10^{-5}$
Lagrange Multiplier ARCH Test	ARCH(5): $X^2 = 132.84$, p-value $< 2.2 \times 10^{-16}$ ARCH(10): $X^2 = 159.23$, p-value $< 2.2 \times 10^{-16}$

Figure 6 confirms the results obtained above. The blue curve represents a normal distribution, while the purple curve represents the distribution of our series. We can see that the distribution of returns is more concentrated in the centre and that the tails of the distributions are thicker. The figure also confirms that V2 is positive, indicating a leptokurtic distribution.

Figure 6: Distribution of returns values

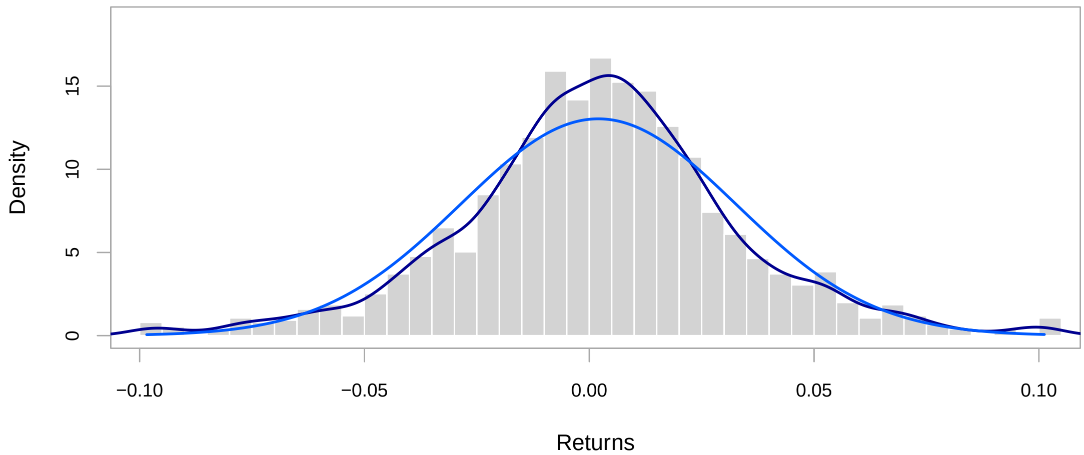

3 Estimation of volatility models over 2018-2022

This second part involves estimating the various models in order to be able to forecast the volatility of returns. To take account of conditinonal heteroscedasticity, we will use models of the ARCH-GARCH class. We will estimate 4 different models: GARCH, GJR-GARCH, IGARCH and Risk-metrics. For each of these four models, we will estimate volatility using two distributions: Normal and Student. This gives us a total of 8 models to estimate. Once the valid models have been identified, we will comment on the best model.

Table II shows us that all the estimated models are correct. In fact, they all respect the constraints of stationarity, positivity and significance of the variables. Although the models seem similar, we note that the GJR-GARCH model stands out from the others with a higher Log-Likelihood and lower information criteria.

Table II: Comparative table of models estimated from a Normal distribution

	Coefficient		$\|t$-value> 1.64	persistance	half-life	log-likehood	Akaike	HQ
GARCH	$\omega > 0$	0.000	3.350	0.969	20.890	3218.560	-4.260	-4.260
	$\alpha \geq 0$	0.099	5.420					
	$\beta \geq 0$	0.870	36.680					
	$\alpha + \beta < 1$	0.969						
GJR-GARCH	$\omega > 0$	0.000	3.410	0.962	17.750	3224.228	-4.267	-4.260
	α	0.058	3.260					
	$\beta \geq 0$	0.860	34.151					
	γ	0.088	2.980					
	$\alpha + \gamma \geq 0$	0.145						
	$\alpha + \beta + (\gamma/2) < 1$	0.962						
IGARCH	$\omega > 0$	0.000	4.879			3212.585	-4.254	-4.250
	$\alpha \geq 0$	0.117	6.526					
	β	0.883						
Riskmetrics	α	0.060				3200.110	-4.240	-4.239
	β	0.940						

The results of the Table II are very close to the previous table since, once again, all the models are significant and the best model is once again the GJR-GARCH model.

Table III: Comparative table of models estimated from a Student distribution

	Coefficient		$\|t$-value> 1.64	persistance	half-life	log-likehood	Akaike	HQ
GARCH	Cst(V) > 0	0.000	3.980	0.979	32.000	3230.562	-4.275	-4.269
	$\alpha \geq 0$	0.100	5.110					
	$\beta \geq 0$	0.880	36.410					
	$\alpha + \beta < 1$	0.980						
	Student	9.380	4.220					
GJR-GARCH	Cst(V) > 0	0.000	2.608	0.974	26.430	3237.819	-4.283	-4.276
	α	0.050	2.770					
	$\beta \geq 0$	0.870	31.137					
	γ	0.110	3.153					
	$\alpha + \gamma \geq 0$	0.160						
	$\alpha + \beta + (\gamma/2) < 1$	0.970						
	Student	9.290	4.861					
IGARCH	Cst(V) > 0	0.000	3.458			3228.638	-4.274	-4.269
	$\alpha \geq 0$	0.110	6.328					
	β	0.890						
	Student	7.830	5.130					
Riskmetrics	α	0.060				3220.400	-4.266	-4.263
	β	0.940						
	Student	8.810	5.611					

We have just seen that for both distributions, the GJR-GARCH model performed better in predicting the volatility of returns. Now, if we compare the two distributions, we can see that the GJR-GARCH student performs better because it maximises the Log-Likelihood and minimises the information criteria. Therefore, we keep this model for the rest of the analysis.

The GJR-GARCH Student model has a persistence of 0.974 (Table III), which means that a volatility shock will disappear over time but it will take time because the persistence is relatively close to 1. As for the half live, it is 26.430 which means that following a volatility shock, it will take about 26 days to return to the mean.

To verify that our GJR-GARCH Student model is valid, we need to test certain hypotheses on the residuals (Table IV). All the tests on the residuals have a p-value greater than 0.05, which means that we do not reject the null hypotheses of non-autocorrelation and homoscedasticity of the residuals. Thus, at the 1% risk threshold, there is no evidence of autocorrelations of the residuals or of homoscdasticity of the residuals. Therefore, the model is valid.

Table IV: Diagnosis of the residuals of the GJR-GARCH model with a Student distribution

Test	**Value**	**p-value**
Ljung-Box Test on Standardized Residuals	$Q(5) = 1.527$	0.7332
Weighted Ljung-Box Test on Standardized Squared Residuals	$Q2(5) = 3.1324$	0.3834
Weighted LM-ARCH(5)	4.394	0.1416

4 Volatility forecast for 2023

In this section, we will forecast the volatility of NVDA asset returns over the period 2023. To this end, we will predict the variance for the models we validated earlier. Once the models have been estimated, we can compare the models in terms of forecast accuracy and accuracy comparison tests.

Table V: Results of the models under the Normal and Student distributions

Normal distribution				
Model	MSE	Rank R	MCS R (pvalue)	R2OOS
GARCH	0.95	6	0.15	0.00042
GJR-GARCH	0.95	5	0.69	0.0013
IGARCH	0.95	1	1.0000	-0.0011
Riskmetrics	0.95	4	0.92	—

Student distribution				
Model	MSE	Rank R	MCS R (pvalue)	R2OOS
GARCH	0.95	8	0.10	-0.0001
GJR-GARCH	0.95	2	1.0000	0.0011
IGARCH	0.95	7	0.15	-0.0011
Riskmetrics	0.95	3	1.0000	0.005

Table V presents the forecasting performance of the estimated models under both normal and Student distributions. The Mean Squared Error (MSE) remains constant across models, indicating that MSE alone is not sufficient to differentiate forecasting quality.

Regarding the out-of-sample R^2 (R2OOS), we observe that the GJR-GARCH model consistently outperforms the Riskmetrics model, regardless of the assumed distribution. This suggests that capturing asymmetry in volatility dynamics contributes to better forecast accuracy. The IGARCH model exhibits negative R2OOS values, indicating poor forecasting performance and a tendency to overfit past volatility patterns.

The Model Confidence Set (MCS) test results confirm that the models have statistically equivalent predictive capacities, as all p-values exceed the 0.05 threshold. This implies that even if there were slight variations in MSE, they would not be statistically significant, reinforcing the robustness of multiple models in predicting NVDA's volatility.

Table VI: Matrix of p-values of Diebold-Mariano tests between models (normal distribution)

Models	GARCH	iGARCH	GJR-GARCH	Riskmetrics
GARCH	-	0,408	1,000	1,000
iGARCH	0,408	-	0,434	1,000
GJR-GARCH	1,000	0,434	-	0,584
Riskmetrics	1,000	1,000	0,584	-

To complete the comparison of the predictive capacities of the models, we can use the Diebold Marianno (DB) Test. Using Tables VI, VII and VIII, we observe that all the p-values are greater than 0.05. At the 5% risk threshold, we cannot reject the null hypothesis. Therefore, we conclude that the models provide the same forecasting quality.

Table VII: Matrix of p-values of Diebold-Mariano tests between models (Student distribution)

Models	GARCH	iGARCH	GJR-GARCH	Riskmetrics
GARCH	-	1	1	0,973
iGARCH	1	-	1,000	0,725
GJR-GARCH	1	1	-	1
RISK	0,973	0,725	1,000	-

The cross-distribution DM test results in Table VIII confirm that assuming a Student distribution instead of a Normal distribution does not lead to significant differences in forecast accuracy. This suggests that heavy-tailed distributions may not provide a notable advantage in predicting NVDA volatility within our sample period.

Table VIII: Matrix of p-values of Diebold-Mariano tests between models (Student vs Normal distributino)

Models	GARCH student	iGARCH student	GJR-GARCH student	Riskmetrucs student
GARCH normal	0,426	0,420	1	1
iGARCH normal	1	1	1	1
GJR normal	0.4882	1	1,000	0,585
RISK normal	0.9743	1	0,581	1

In conclusion, our analysis indicates that while some models (such as GJR-GARCH) slightly outperform others in certain metrics, overall, all models provide statistically equivalent forecasting accuracy.

5 Conclusion

This study investigates the volatility of Nvidia's stock returns over the period 2018-2023 using GARCH-family models. Our analysis began with an examination of Nvidia's financial and market characteristics, highlighting its dominant position in the semiconductor industry and its strong influence in AI-related markets. We then explored the statistical properties of NVDA returns. To model this volatility, we estimated four different GARCH-family models: GARCH, GJR-GARCH, IGARCH, and RiskMetrics, under both normal and Student distributions. Our results showed that all models satisfy stationarity and significance constraints, with the GJR-GARCH model consistently standing out as the most effective in capturing asymmetric volatility shocks. However, while the GJR-GARCH model demonstrated slightly better predictive performance, statistical comparison tests, including the Diebold-Mariano test and the Model Confidence Set, indicated that no single model significantly outperformed the others in forecasting accuracy. As Nvidia continues to dominate the AI and semiconductor industry, its stock remains highly reactive to market trends and technological advancements. The increasing demand for computing power, coupled with macroeconomic factors, will likely sustain the stock's volatility in the coming years.