MASTER THESIS

Stochastic projections as a risk management and valuation tool applied to Novartis

by

Cédric Masset

Supervised by

Pierre Guillemin

Master of Science in Finance

Faculty of Business and Economics

UNIL | Université de Lausanne

2024

Abstract

Ce mémoire de master explore l'utilisation des projections stochastiques en tant qu'outil de gestion des risques et d'évaluation financière, avec une application spécifique à Novartis. Il analyse de manière critique l'efficacité des méthodes DCF enrichies par des projections stochastiques des FCFF, ainsi que la dynamique des rendements boursiers à l'aide de modèles stochastiques. Ces deux approches sont comparées aux moyennes des estimations des analystes qui utilisent des modèles DCF basés sur des prix cibles ponctuels. En utilisant des données historiques pour modéliser les futurs prix des actions, ce travail prend Novartis comme étude de cas pour déterminer si la médiane des prix obtenus par les deux types de projections stochastiques offre des prévisions plus ou moins précises que la moyenne des prix cibles proposés par les analystes. À travers une analyse quantitative, cette étude vise à fournir une compréhension nuancée de la fiabilité prédictive de ces méthodes d'évaluation financière.

This Master's thesis explores the application of stochastic projections as a risk management and financial valuation tool, with a specific application to Novartis. It critically analyses the efficiency of DCF methods enhanced by stochastic projections of FCFF, as well as the dynamics of stock returns using stochastic models. These two approaches are compared with the average estimates of analysts who use DCF models based on point target prices. Using historical data to model future stock prices, this work takes Novartis as a case study to determine if median prices obtained by the two types of stochastic projections offer more or less accurate forecasts than the average target prices proposed by analysts. Through quantitative analysis, this study aims to provide a nuanced understanding of the predictive reliability of these financial valuation methods.

Contents

List of Figures

Abbreviations

NVS Novartis

IQR **I**nter**q**uartile **R**ange

MAD **M**edian **A**bsolute **D**eviation

DCF **D**iscounted **C**ash **F**low

FCFF **F**ree **C**ash **F**low to the **F**irm

MCS **M**onte-**C**arlo **S**imulation

WACC **W**eighted **A**verage **C**ost of **C**apital

CAPM **C**apital **A**sset **P**ricing **M**odel

AIC **A**kaike **I**nformation **C**riterion

TGR **T**erminal **G**rowth **R**ate

AR **A**uto **R**egressive

GARCH **G**eneralized **A**uto**R**egressive **C**onditional **H**eteroskedasticity

ACF **A**uto**C**orrelation **F**onction

PACF **P**artial **A**uto**C**orrelation **F**onction

VaR **V**alue **a**t **R**isk

CVaR **C**onditionnal **V**alue **a**t **R**isk

PDF **P**robability **D**ensity **F**onction

1 Introduction

1.1 Analysis objective

Research objective

The main objective of this research is to examine and compare two approaches to the financial valuation of equities on the stock markets. We will analyse the results obtained using the traditional Discounted Cash Flow (DCF) method, which is commonly used by financial analysts to incorporate forecasts of fundamental elements such as expenses, income and capital expenditure. These results will be compared with those from two methods based on stochastic projections, which are based on statistical assumptions rather than on specific judgements about the company's future performance. The effectiveness of these methods will be assessed by comparing them with actual market share prices, providing a direct measure of their predictive accuracy. We will use Novartis as the basis of analysis for this research.

More specifically, the aim of this work is to determine whether analyst consensus price estimates, using DCF models, provide more accurate or less accurate predictions of the true share price compared to methods that use stochastic Free Cash Flows to Firm (FCFF) projections embedded in a DCF framework to provide valuation. In addition, we will analyse the dynamics of stock returns using historical price data to project future share prices. This will allow us to juxtapose an approach based on historical fundamental data (FCFF) with another that focuses on historical share price dynamics.

Research question

The research question is: "Are median stochastic projections, as a pricing method, more accurate than analyst price valuations performed by traditional DCFs?" This question guides our exploration of the various valuation strategies and aims to identify which of these methodologies offers the greatest accuracy and reliability in forecasting share prices.

2 FCFF Projections

2.1 FCFF and valuation

Free Cash Flow to the Firm (FCFF) is a crucial financial indicator for assessing a company's intrinsic value. FCFF represents the amount of cash a company can generate after covering its operating expenses and capital investments, but before meeting its financial charges such as interest payments. This measure provides a clear picture of the company's ability to generate free cash flow to repay creditors, pay dividends and reinvest in its operations without having to resort to new debt or external financing.

To calculate FCFF, we generally start with earnings before interest, tax, depreciation and amortisation (EBITDA), then subtract the tax that would have been paid if the company had no debt (EBIT(1-Tax rate)). From this, we then subtract capital investment (capital expenditure) and add changes in working capital requirements, to obtain the free cash flow available to all the company's suppliers of capital.

The use of FCFF is particularly relevant in the Discounted Cash Flow (DCF) valuation method. This approach is based on the idea that the value of a company is essentially the present sum of all its future cash flows available to repay investors after financing operations and growth. The company's future cash flows are estimated for a typical projection period (often 5 to 10 years) and then discounted at a rate that reflects the risk associated with those cash flows. This rate is generally the company's weighted average cost of capital (WACC), which incorporates the cost of debt and the cost of equity according to their respective proportions in the company's financing structure.

The discounting process converts future cash flows into a net present value (NPV), providing an estimate of the total value of the business, including its debt. To arrive at the equity value, net debt (total debt less cash equivalents and short-term investments) is subtracted from the enterprise value obtained. This equity value is particularly useful for shareholders because it allows them to compare the current market value of the company's shares with the value calculated from the FCFF, to judge whether the shares are overvalued, undervalued, or correctly valued.

The main advantage of using FCFF in DCF models is that they are not affected by the company's capital structure. This provides a purer measure of the company's economic performance, independent of how it is financed. However, this method assumes an ability to accurately estimate future cash flows and the appropriate discount rate, a task often complicated by the uncertainty of economic forecasts and market fluctuations.

The FCFF formula is as follows:

$$FCFF = \text{EBIT} - \text{Taxes} + \text{Depreciation} - \text{CAPEX} - \Delta\text{NWC}$$

2.2 Polynomial Regressions on FCFF

We applied degree 1-3 polynomial regressions to project FCFF over a ten-year period. The image clearly shows this analysis, with historical data for Novartis FCFF from 1996 to 2023, followed by projections to 2033 based on different regression models.

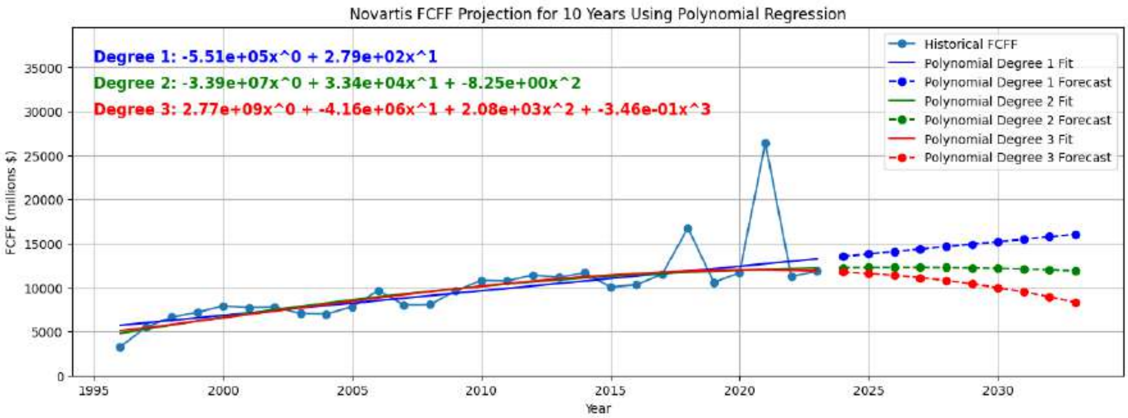

Figure 1: Polynomial Regressions

First-degree polynomial regression, or linear regression, is the simplest and most straightforward form of analysis. In this model, the FCFF is predicted from a linear relationship with time, suggesting constant growth over the years. This method is often used for its simplified approach and ease of interpretation, as shown by the blue line on the graph, which offers a general view of the direction cash flows could take, assuming that past conditions continue without major changes.

By switching to a second-degree polynomial regression, we introduce a quadratic term that captures curvature effects in the historical data. This can reflect periods of accelerated growth or slowdown, providing a more nuanced picture of the dynamics of a time series.

Third-degree regression adds a further cubic term, allowing even greater flexibility in modelling a time series. This model, represented by the red curve in the image, can theoretically adjust for complex patterns such as oscillations or trend reversals that might occur as a result of significant changes in the company's economic, regulatory or market environment. However, while this model offers increased adaptability, it can also be prone to problems of overfitting, where the model conforms too closely to historical data, compromising its ability to correctly predict future trends.

However, it is crucial to note that increasing the degree of polynomial regression does not necessarily imply better prediction. Indeed, although higher degree models may appear to fit historical data better, they may also become overly sensitive to minor fluctuations or anomalies in the data, leading to forecasts that appear unlikely or highly volatile for future years. This is particularly evident in the case of third-degree forecasts, where projections decrease exponentially.

In sum, although polynomial regressions offer a useful tool for sketching future trends based on past performance, it is essential to choose the degree of the regression with caution. A model that is too simple may not capture the full complexity of the data, while one that is too complex may lead to unrealistic predictions. The image shows this dilemma, with FCFF projections becoming increasingly unstable and unrealistic as the degree of regression increases.

2.3 AR Model on FCFF

2.3.1 Stationnarity test of FCFF

Another technique for projecting FCFF is to parameterise an AR model. However, before parameterising an AR model, it is necessary to test the stationarity of the time series. The ADF (Augmented Dickey-Fuller) test was used to assess the stationarity of the FCFF. The results show an ADF statistic of 0.686 and a p-value of 0.989, well above the 0.05 threshold. The critical values at 1%, 5% and 10% are -3.724, -2.986 and -2.633 respectively. Since the p-value is very high and the ADF statistic is above the critical values, the time series is non-stationary. This implies that this non-stationarity may pose problems for the parameterisation of the AR model, as AR models generally require stationary series to produce reliable results.

2.3.2 ACF and PACF

The graph below shows the autocorrelation function (ACF) and partial autocorrelation function (PACF) of the FCFF. These tools are used to identify the structure of time dependence in the data and to determine the appropriate order of time series models, such as AR (Autoregressive) models.

On the left-hand graph, we have the ACF, which shows the correlation of current

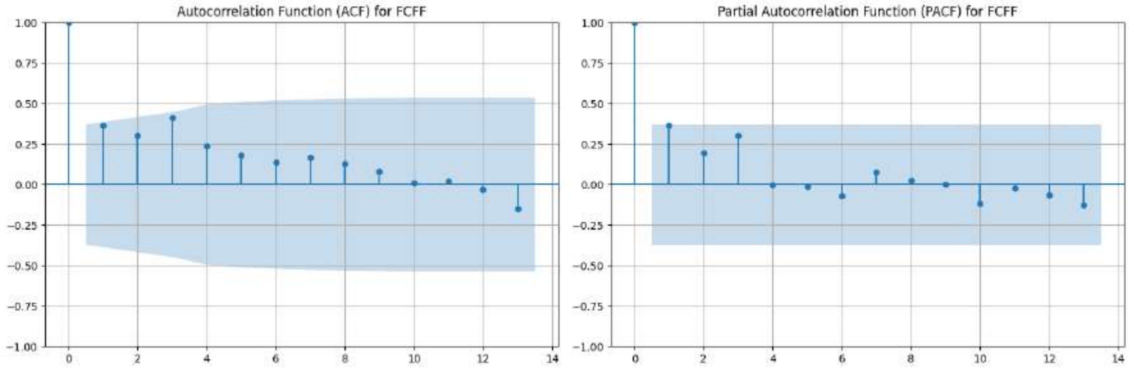

Figure 2: ACF and PACF on historical FCFFs

observations with the values at different lags. We can see that the values at lags 1 to 3 are relatively higher than the others, suggesting a certain correlation between the current ACF values and those of the previous period. However, these correlations at lags 1 to 3 are not excessively strong. The other blue bars representing the other correlations are also present, but less significant.

The graph on the right represents the PACF, which measures the correlation of observations with past lags, after eliminating the effects of intermediate lags. Here too, lag 1 appears to be more significant than the others, slightly exceeding the limits of the confidence interval, while the subsequent lags remain broadly within the limits, indicating that their direct influence is less after taking lag 1 into account.

This analysis suggests that the FCFF data show a slightly more noticeable dependence on lag 1, although this dependence is not extremely pronounced. Therefore, an AR(1) model seems to be a reasonable choice to capture the dynamics of the data, as it takes into account the relationship between the current FCFF values and the values of the previous period. However, given that lag 1 is not clearly significant, it may be beneficial to also explore other models or refine the AR(1) model to ensure that it best represents the characteristics of the time series.

2.3.3 AR model settings

The formula for the first-order AR model for FCFF is given by :

$$\text{FCFF}_t = \alpha + \phi \text{FCFF}_{t-1} + \varepsilon_t$$

where $\alpha = 6556.7892$ (p-value $= 0.000$) and $\phi = 0.3705$ (p-value $= 0.028$). ε_t represents the error term at year t.

We applied the AR(1) model to Novartis' historical FCFF to assess its ability to predict future cash flows. For comparison, a second-order AR model (AR(2)) was also used. The image shows the predictions of both models against the actual FCFF values. The blue dots represent the actual FCFF values, while the orange and green dotted lines represent the predictions of the AR(1) and AR(2) models respectively.

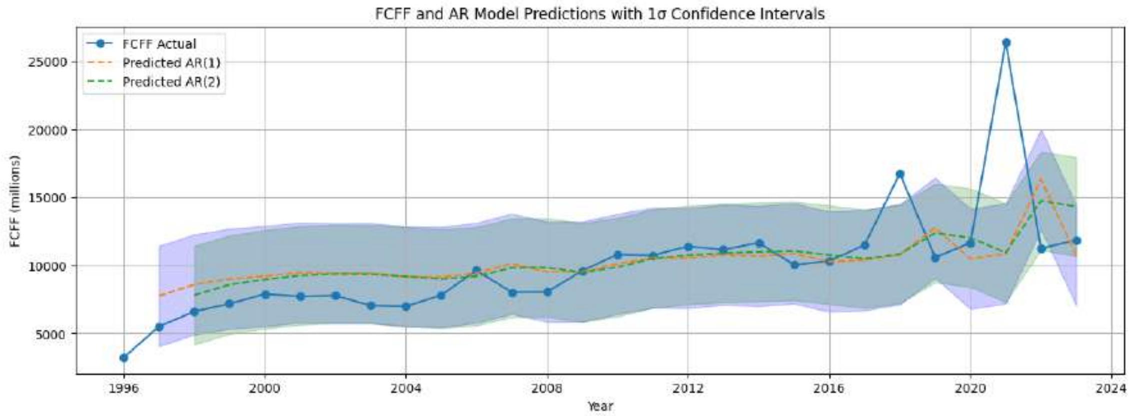

Figure 3: Application of AR models

The confidence intervals, illustrated by the shaded areas around the predictions, correspond to one standard deviation for the residuals, providing an indication of the uncertainty associated with the predictions. The confidence interval for AR(1) is shown in blue, while that for AR(2) is in green. These intervals show the expected variability around the predictions, indicating where the actual values of the FCFF could lie with a certain level of confidence.

Comparing the two models, we can see that the AR(1) model offers predictions that closely follow the general trend of the historical data, but with certain limitations in terms of accuracy. The AR(2) model, while offering a slight improvement in some cases, also shows similar challenges. Both models capture broad trends but may have difficulty predicting the steeper variations observed in real data. These analyses provide a better understanding of the strengths and limitations of AR models for forecasting corporate cash flows.

2.3.4 Tests on the residuals of the AR model

Analysis of the residuals from the AR(1) model fitted to the data shows that they do not follow a normal distribution. After fitting the regression model to the FCFF data, the residuals were extracted to perform a normality test.

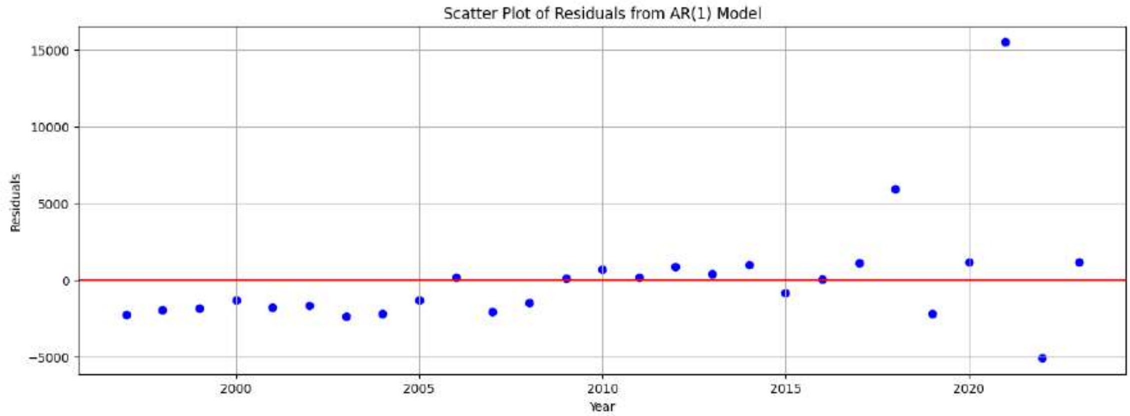

Figure 4: Residuals of AR(1) Model

The test used to examine the normality of the residuals is the Shapiro-Wilk test, which evaluates the null hypothesis that the residuals are normally distributed. The test result shows a very high statistic and an extremely low p-value ($p < 0.05$), clearly indicating that the normality hypothesis must be rejected. This means that the distribution of the residuals deviates significantly from a normal distribution, which could affect the validity of certain statistical tests or confidence intervals based on the assumption of normality.

Analysis of the ACF and PACF graphs of the AR model residuals (below) shows that the AR model has succeeded in reducing the autocorrelation present in the initial FCFF time series, although some lags show variations.

In the graphs, we observe that the first two lags of the residuals have a lower autocorrelation compared to the initial time series. This indicates that the AR model has succeeded in capturing and effectively modelling short-term temporal dependencies.

However, the third lag of residuals shows a slightly higher autocorrelation than that observed in the initial time series. This increase in autocorrelation at the third lag could suggest that the AR model has not completely eliminated all temporal dependencies or that there are structures in the data that the AR model has not fully captured.

In summary, although the AR model has reduced the autocorrelation of the first two lags, the third lag shows a higher autocorrelation than in the original time series.

This indicates that, despite an overall improvement, some temporal dependencies remain, perhaps requiring further investigation or additional adjustments to the model.

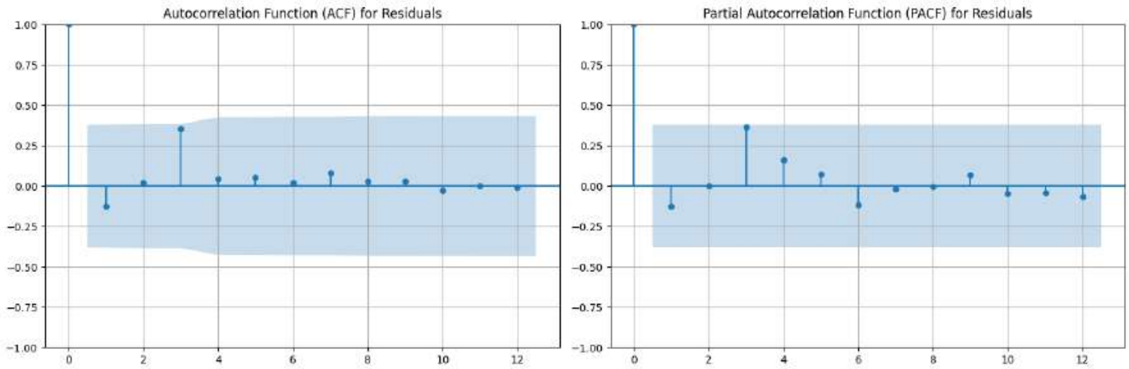

Figure 5: Autocorrelation of AR model residuals

2.3.5 Conclusion

In conclusion, the use of the AR model to predict Novartis FCFF presents mixed results. Although the ACF and PACF analysis showed that the AR model reduced some of the autocorrelations present in the initial time series, particularly for the first two lags, the autocorrelation at the third lag of the residuals is higher than in the initial time series. This suggests that the AR model has not fully captured all the temporal dependencies.

Furthermore, although autocorrelation is present in the data, it does not seem significant enough to fully justify the use of an AR model. This low significance of autocorrelation limits the effectiveness of the AR model for this time series.

Another crucial point is that the residuals of the AR model do not follow a normal distribution. This non-normality of the residuals is problematic because it prevents a correct simulation of the forecast uncertainty of the AR model. A non-normal distribution of residuals compromises the validity of confidence intervals and statistical tests based on the normality assumption.

In sum, the results obtained indicate that the AR model does not provide sufficiently robust and reliable forecasts for the Novartis FCFF, due to the low significance of the autocorrelation and the non-normality of the residuals. It would be appropriate to explore other models or approaches to improve the accuracy of forecasts and the representation of uncertainty.

2.4 Projections by distribution of changes in FCFF

Rather than using an econometric model to project the FCFF with residuals that fail to be normally distributed, we have chosen to model variations in the FCFF directly from one year to the next by fitting a Student distribution (t-Student). The t-Student distribution allows us to better capture the central concentration with a very low degree of freedom, unlike the normal distribution. The following image illustrates this approach, showing the time series of percentage variations in the Novartis FCFF and the histogram of these variations with the t-Student fit.

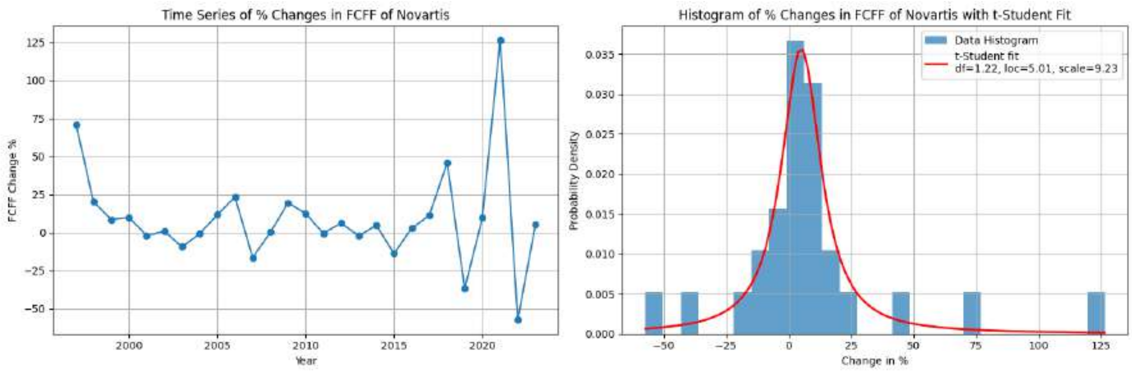

Figure 6: Distribution of FCFF variations

On the basis of the parameters of the Student distribution (t-Student) adjusted for variations in the FCFF, we can make random projections of future variations using this distribution. The formula used for these projections is as follows: $\text{FCFF}_t = \text{FCFF}_{t-1} \times (1 + \epsilon_t)$, where ϵ_t follows a t-Student distribution with parameters $\mu = 0.0501$, $\sigma = 0.0923$, and $df = 1.22$.

Applying this formula, we generate (in the image below) several time series of projected FCFF for the period 2024 to 2033, allowing us to visualise different possible trajectories of Novartis FCFF. The image shows the results of a 10-year Monte Carlo simulation based on t-student adjusted FCFF variations.

The blue dots represent the historical series of FCFF from 1996 to 2023. The dark blue lines represent various individual trajectories randomly generated during the Monte Carlo simulation. The red line illustrates the median FCFF projections, while the green and orange lines show the 25th and 75th percentiles of the projections respectively, indicating the pessimistic and optimistic scenarios among the projections.

The aim of this simulation is to produce several time series of projected FCFF for the years 2024 to 2033. These time series will then be discounted to one year (December 2024 or January 2025) using the discounted cash flow (DCF) method.

This approach better captures the uncertainty and potential variability of DCFs, thus providing a range of possible values for the Novartis financial valuation.

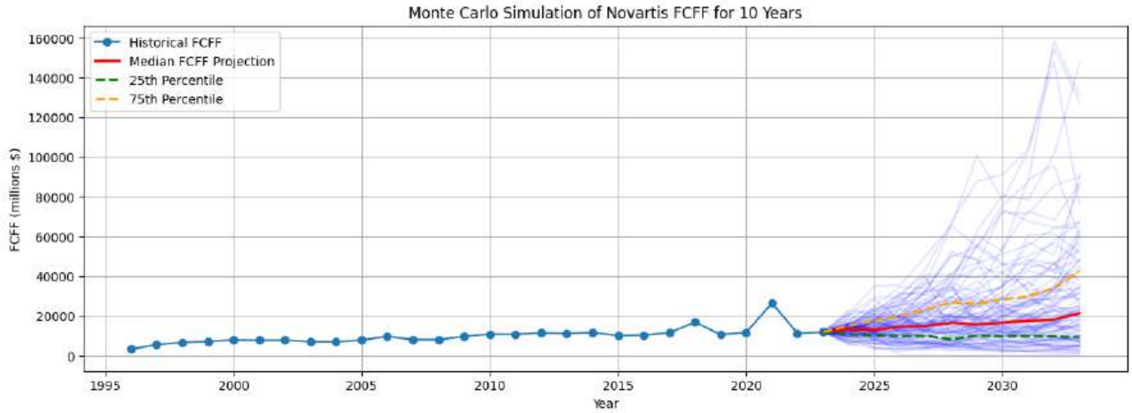

Figure 7: Projections of FCFF

2.5 Discount rate - WACC

To discount our projected FCFF, we will need an appropriate discount rate. In order to do this, we need to calculate the WACC (Weighted Average Cost of Capital).

2.5.1 Beta

The graph shows a linear regression of the logarithmic daily returns of Novartis (NVS) against the logarithmic daily returns of the Swiss Market Index (SMI). The regression line is used to determine the Beta of Novartis.

To begin with, we calculated the daily total log-returns for Novartis and the SMI. The log-return is given by the formula:

$$\text{Total Log-Return} = \ln\left(\frac{\text{Adj Close}_t}{\text{Adj Close}_{t-1}}\right)$$

where Adj Close_t is the closing adjusted price at time t.

We then performed a linear regression of Novartis returns (y-axis) against SMI returns (x-axis). The linear regression gives us a straight line of the form:

$$y = \alpha + \beta x$$

where y is the Novartis return, x is the SMI return, α is the intercept and β is the slope of the line.

The slope of the regression line, indicated by β, represents the Beta of Novartis. Beta measures the sensitivity of Novartis returns to market returns represented by the SMI. In the graph, the regression line is given by:

$$y = 0.0002 + 0.6123x$$

Here, $\beta = 0.6123$. This means that for every 1% change in the SMI, the Novartis return changes by an average of 0.6123%. A Beta of 0.6123 indicates that Novartis is less volatile than the market (SMI). If $\beta > 1$, this would mean that Novartis is more volatile than the market. If $\beta < 1$, as is the case here, Novartis is less volatile.

The red line on the graph represents the regression line, and the equation for the line is $y = 0.0002 + 0.6123x$. The blue dots represent the actual daily returns of Novartis compared to the daily returns of the SMI.

In conclusion, we have found the Beta of Novartis from the logarithmic daily returns using simple linear regression, and the Beta is 0.6123, which will be used in the calculation of the WACC.

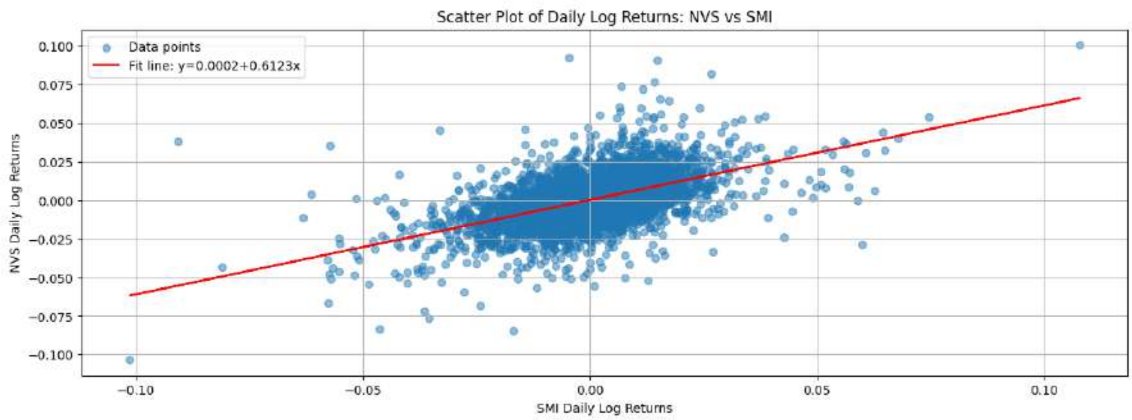

Figure 8: Novartis Beta

2.5.2 WACC

This table shows the calculation of the WACC of Novartis, using different financial components.

Table 1: WACC

Cost of Equity (CAPM)	
10-year risk-free rate (R_f)	0,040
Levered Beta (β)	0,612
Equity Risk premium (ERP)	0,055
Country Risk Premum (CRP)	0,000
Cost of Equity (R_e)	0,074
% of Equity	0,791
Cost of Debt	
SP Global Ratings of Novartis	AA
Cost of Debt	0,052
% of Debt	0,209
Corporate Tax Rate	0,180
WACC	0,067

The risk-free rate used, 4%, comes from the Federal Reserve Bank of Saint-Louis[1] for 10-year bonds. The Levered Beta calculated is 0.612. The Equity Risk Premium and Country Risk Premium are taken from the work of Aswath Damodaran[2]. The Cost of Equity, calculated using the CAPM, is 0.074, representing the return expected by shareholders. The proportion of equity financing is 79.1%. The financial strength of Novartis is recognised with an 'AA' rating from Standard & Poor's[3]. The cost of debt of 0.052 is based on current market conditions and Novartis financial documents. The capital structure includes 20.9% debt. The corporate tax rate is 18%. This information provides a basis for estimating the WACC of Novartis at 6.7%. All fundamental values are calculated directly from Novartis financial documents.

[1] https://fred.stlouisfed.org
[2] https://www.stern.nyu.edu
[3] https://www.novartis.com

2.6 Monte-Carlo simulation on DCF

2.6.1 Principle and assumptions

To estimate the share price using the DCF model, we apply a discounting method starting from the year 2025. In our analysis, the FCFF for the years 2025 to 2033 are the only ones taken into account for discounting. The FCFF for 2024 is not included, as it represents the current reference year, and therefore does not need to be discounted for a one-year assessment.

By incorporating these discounted FCFFs into a Monte Carlo simulation (including a terminal value), we can dynamically assess the impacts of these future flows. This method makes it possible to address uncertainties and model various potentially impactful scenarios. By discounting each FCFF from 2025 to 2033 to reflect its value in 2025 (late 2024 or early 2025), we obtain enterprise values.

Once we have obtained these discounted cash flows, we divide them by the total number of shares in the company to obtain directly the expected share price for the year 2025. This division transforms the sum of the discounted values into an estimate of the price per share, providing a direct and relevant measure of value for investors.

The formula for the price of a share in one year's time by discounting FCFF is given by :

$$P_1 = \frac{\sum_{t=2}^{n+1} \frac{FCFF_t}{(1+r)^{t-1}} + \frac{\frac{FCFF_{n+1} \times (1+g)}{r-g}}{(1+r)^n}}{N}$$

where :

- $\sum_{t=2}^{n+1} \frac{FCFF_t}{(1+r)^{t-1}}$ represents the discounted sum of the FCFF, which we discount to next year.
- g is the terminal growth rate.
- r is the discount rate (WACC)
- N is the total number of shares.
- n : Number of years of FCFF projections.

2.6.2 Applying Monte Carlo simulation to FCFF

The parameters presented in the table below summarise those used in the Monte-Carlo simulation. They concern both the simulation of variations in FCFF (parameters of the t-Student law), and the parameters for updating the simulated FCFF.

Table 2: Monte-Carlo DCF Parameters

Parameters	Values
WACC	6.70%
TGR	2.00%
Historical Mean of variation	5.01%
Historical Volatility of variation	9.23%
Degree of freedom of variations	1.22
Inital FCFF (mio $)	11864
Number of shares	2044000000
Number of projected years	10
Number of simulations	10000

2.6.3 One-year price distribution

The probability density function (PDF) of the final prices was modelled using an asymmetric t-distribution, as shown in the image below. This choice of distribution is significant because it takes into account the asymmetry of possible outcomes. The asymmetry suggests that there is a non-negligible probability of above-average results, which corresponds to the fact that the average price is higher than the median.

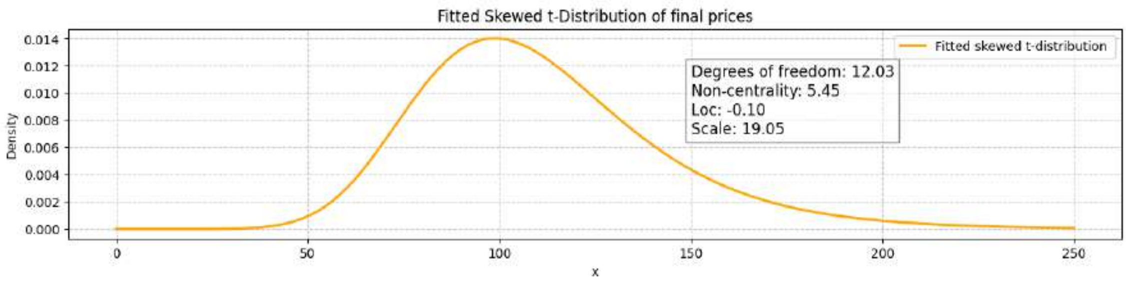

Figure 9: PDF of final prices (based on FCFF)

This asymmetry can be interpreted as market sentiment indicating that, despite potential risks, there is optimism for significant upward movements, probably due to favourable market or company-specific catalysts. The degree of freedom and non-centrality parameters indicate the level of kurtosis and skewness respectively, providing information on the tail risk - the probability of extreme outcomes.

The results of the Monte-Carlo simulation of FCFF indicate that the projected median share price is \$106.62. This median value is a crucial point of comparison with the average of \$111.04, highlighting a possible asymmetry in the distribution of possible outcomes. Given that the current share price at the end of 2023 is \$100.97, the projected median implies a potential return of 5.6% over the year.

3 Stochastic projection of returns

3.1 Choice of sample period

When we project returns, we define their dynamics in terms of observed empirical characteristics. However, it is essential to choose the length of the sample period carefully. Opting for a longer period means that the dynamics of the stock faithfully reflect the entire period under consideration. If the stock shows a particularly marked trend and we wish to accentuate this trend in the stochastic projections, it may be wise to select a shorter sample period. This avoids diluting recent strong directional returns with an average. The choice of sample length is therefore crucial and should be adapted to give a more specific drift, depending on analytical needs.

Continuing with this logic of selecting the sample period for return projections, let's look at two hypothetical examples of companies that illustrate the importance of this approach: a mature company and a very fast-growing company.

Let's start with a mature company. For such a company, share price fluctuations may be relatively stable and predictable, reflecting moderate growth and regular revenues. In this case, choosing a long sample period for return projections is appropriate. This allows a full range of economic cycles and market fluctuations to be incorporated, providing a holistic view of company performance that helps stabilise projections by smoothing out short-term anomalies. In contrast, consider a company in the technology sector, characterised by very strong growth. For such a company, shares may show extremely volatile price fluctuations due to rapid innovation, regulatory changes, or market reactions to new product announcements. In this context, using a shorter sample for stochastic projections may be more relevant. This allows recent and more relevant trends to be captured without being obscured by historical data that may no longer reflect the company's current situation. In this way, projections place greater emphasis on recent growth dynamics and are potentially more accurate in predicting future returns in a fast-changing environment.

In summary, the selection of the sample period is crucial and must be adapted to the specific context of each company. It has a direct impact on the accuracy and relevance of return projections, particularly in scenarios where companies go through very divergent growth phases.

For Novartis, given that the company is considered to be in a mature phase, we opt for a period that encompasses the main economic crises in order to incorporate these events into the probabilities of our analyses. The period selected runs from 1997 to 2023 inclusive.

3.2 Statistical properties

The following table shows the statistical properties of Novartis' logarithmic returns from 1997 to 2023. The statistics are calculated at three different frequencies: daily, monthly and annually.

Table 3: Summary statistics of log-returns for Novartis (NVS)

Statistic	Close	Adjusted Close
Daily frequency (6,793 observations)		
Mean (%)	0.020	0.031
Median (%)	0.025	0.035
Standard Deviation (%)	1.355	1.346
MAD (%)	0.705	0.706
IQR (%)	1.411	1.410
Ann. mean (%)	5.091	7.767
Ann. std. dev. (%)	21.512	21.369
Skewness	0.025	0.064
Kurtosis	7.555	7.654
Monthly frequency (323 observations)		
Mean (%)	0.422	0.646
Median (%)	0.823	1.089
Standard Deviation (%)	5.210	5.213
MAD (%)	3.675	3.459
IQR (%)	7.370	7.265
Ann. mean (%)	5.067	7.746
Ann. std. dev. (%)	18.049	18.059
Skewness	-0.329	-0.324
Kurtosis	3.493	3.513
Annual frequency (26 observations)		
Mean (%)	3.918	6.692
Median (%)	5.872	9.211
Standard Deviation (%)	13.554	13.885
MAD (%)	8.642	8.446
IQR (%)	18.301	19.309
Skewness	-0.587	-0.709
Kurtosis	2.817	3.128

3.2.1 Analysis of returns

In analysing the statistical properties of Novartis' logarithmic returns between 1997 and 2023, it is crucial to distinguish between the results obtained from "Close" and "Adjusted Close" prices. The differences observed in the two data series reflect adjustments for corporate events such as dividends, etc., which can significantly influence the calculated returns.

For the daily data, we observe that the average return for the adjusted price (0.031%) is slightly higher than that for the closing price (0.020%). The calculation of the median also shows a slight increase for the adjusted price (0.035% versus 0.025%), suggesting a positive impact from the reinvestment of dividends.

In terms of the dispersion of returns, measured by the standard deviation and MAD (Mean Absolute Deviation), the values are very close between the two sets of data, indicating that the adjustments do not substantially affect the overall volatility of returns over the 3 frequency types. This means that the type of price has an implication on the drift of the time series and not on the volatility.

When we look at the annual data, the difference becomes more marked. The mean of adjusted returns (6,692%) is significantly higher than that of unadjusted returns (3,918%), with a standard deviation that is also higher (13,885% versus 13,554%). This shows that dividends have a significant impact on returns measured over long periods.

As far as skewness and kurtosis are concerned, the differences between closing and adjusted prices are not very pronounced, indicating that the distribution patterns of returns remain relatively similar despite the adjustments. This further reinforces the finding of the impact on drift rather than their fundamental distribution.

In conclusion, the statistics show that although adjusted and unadjusted prices follow similar patterns in terms of distribution, the adjustments affect measures of central tendency, particularly over longer periods. This highlights the importance of correctly choosing the time series to use in order to utilise the correct statistical properties.

3.3 Types of Price

We saw earlier that it is essential to choose the right type of price when modelling the future path of a company's share price. Indeed, the choice of price data directly influences the accuracy and relevance of our projections. For our analysis, we focus on unadjusted prices ('Close').

Our aim is to model the dynamics specific to the share price, without incorporating the additional returns linked to reinvested dividends. By using these unadjusted prices, we can capture only the real variations in the share price, giving us a clearer and more accurate picture of its evolution in the market.

Adjusted prices, on the other hand, are particularly useful when we want to calculate overall performance and compare a company's performance with that of other companies or with benchmark indices. Adjusted share prices make it possible to measure total return, which is relevant in comparative contexts. However, it is not this overall dynamic that interests us here. We seek to isolate and analyse the pure trajectory of the share price without the influence of dividends, in order to better understand its intrinsic evolution on the market.

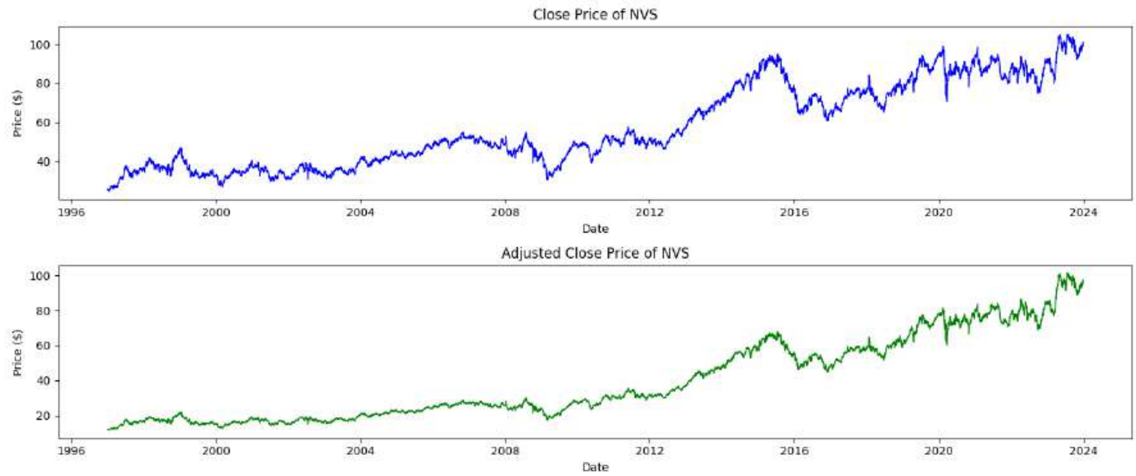

Figure 10: Type of Price

3.4 Normality Test - Jarque-Bera Test

We perform a normality test on Novartis returns to assess whether the data follows a normal distribution. Financial returns are rarely normal, and this check is crucial for our modelling. Indeed, if the returns are not normal, we will have to adjust our analysis methods and choose more suitable models to take account of this non-normality and improve the accuracy of our forecasts.

The following table presents this analysis of the normality of logarithmic returns for Novartis. The normality tests, carried out at various frequencies (daily, monthly and annually), include the calculation of skewness, kurtosis and the results of the Jarque-Bera (JB) test, together with the associated p-values. These statistical indicators are crucial for assessing the distribution of returns on these companies' shares over different periods.

The JB statistic is calculated using the following formula:

$$JB = n\left(\frac{S^2}{6} + \frac{(K-3)^2}{24}\right)$$

where n is the number of observations, S is the skewness, and K is the kurtosis.

Table 4: Normality tests for log-returns of NVS

	Value
	Daily frequency (6,793 observations)
Skewness	0.025
Kurtosis	7.555
JB	5872.937
P-Value	0.000
	Monthly frequency (323 observations)
Skewness	-0.329
Kurtosis	3.493
JB	9.093
P-Value	0.011
	Annual frequency (26 observations)
Skewness	-0.587
Kurtosis	2.817
JB	1.531
P-Value	0.465

The p-value for the Jarque-Bera test is calculated using the chi-square distribution with two degrees of freedom. For a test statistic JB, the p-value is given by the following formula:

$$\text{p-value} = P(X_2^2 > JB)$$

where X_2^2 is a variable which follows a chi-square distribution with 2 degrees of freedom.

This formula is used to determine the probability of observing a value of the JB statistic as extreme or more extreme under the null hypothesis that the data follow a normal distribution.

3.4.1 QQ-Plot

This QQ-plot shows the distribution of Novartis daily returns compared to a theoretical normal distribution. The red line represents an ideal normal distribution. It can be seen that the data broadly follows this line, indicating a similarity to a normal distribution, but with notable deviations. The points at the extremes, particularly for the lower and upper quantiles, move away from the red line, suggesting thicker tails than normal. This implies a more frequent occurrence of extreme returns than predicted by the normal distribution, typical of financial data with risks of unusual losses or gains. The deviations observed in the QQ-plot, particularly the thicker tails, confirm the results of the Jarque-Bera test carried out earlier.

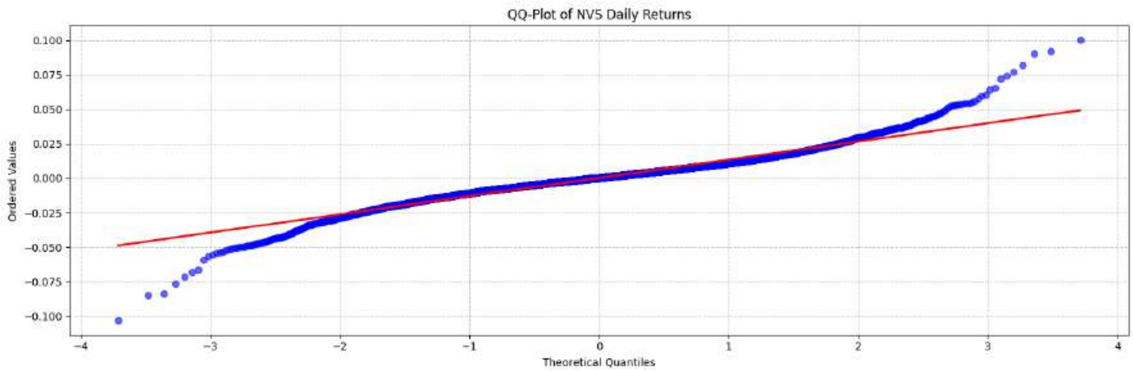

Figure 11: QQ-Plot on NVS

3.5 Fitting distributions on returns

In order to take into account the various characteristics of the returns observed above, we need to choose a distribution that takes into account both the fat tails and the asymmetry of the returns. We will therefore adapt a normal distribution, a t-student distribution and a skewed-t distribution to the returns in order to see which models the returns as closely as possible.

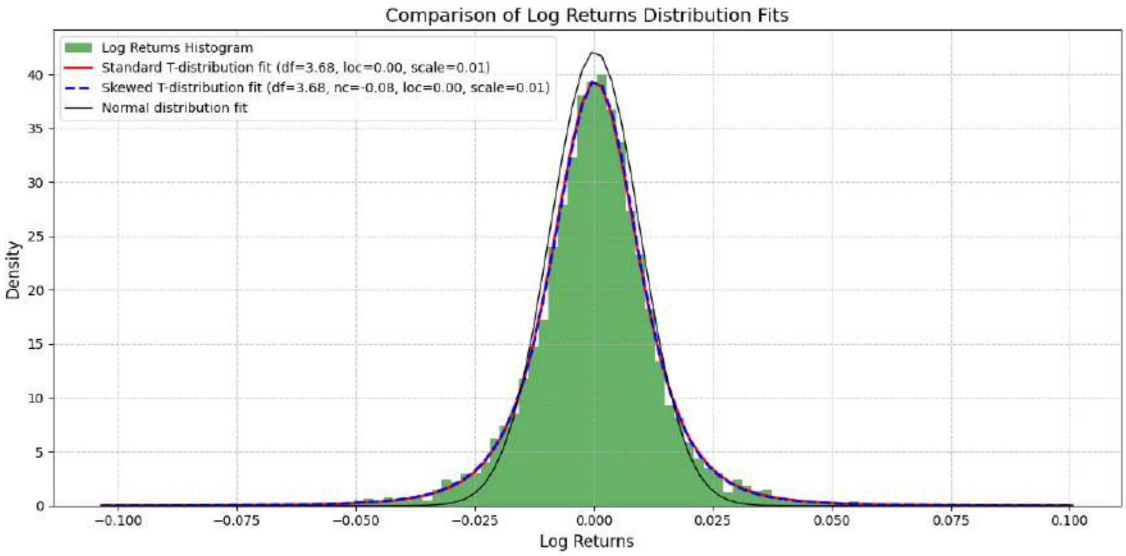

Figure 12: Distribution Fits on NVS

The Akaike Information Criterion (AIC) allows us to determine which distribution is best suited to the returns.

We find that the Student's t distribution and the skewed t distribution come out on top, with very little difference, with the skewed t distribution being slightly better. This is due to the fact that the positive skewness of daily returns is very low (0.025). The lower the AIC value, the better the fit of the distribution to the returns.

Table 5: Statistics of distribution fits

Akaike Information Criterion	**Value**
Normal	-39156.18
t-Student	-40136.55
t-Skewed	-40136.68

The parameters calculated for the t-skewed distribution (adapted to the period 1997 to 2023) correspond to $df = 3.68$, $nc = -0.08$, $\mu = 0.00$ et $\sigma = 0.01$.

3.5.1 Observation on nc parameter

We can see that the non-centrality parameter *nc* in a t-skewed distribution does not correspond directly to the skewness calculated on historical returns (0.025). This parameter does not measure asymmetry in the same way as skewness. Indeed, *nc* adjusts the distribution to better capture data characteristics such as skewness and heavy tails, but does not serve solely to measure skewness as does the skewness coefficient. The skewness calculation is a descriptive measure that assesses the asymmetry of the data distribution around its mean, indicating whether the tails are longer on one side or the other of the mean. In contrast, the *nc* parameter in an adjusted distribution is used to optimise this fit, taking into account not only the skewness but also other aspects of the data to minimise the error between the theoretical distribution and the observations. So, although the two measures address data skewness, they do so in different ways and with different objectives, which is why they do not necessarily coincide.

3.6 Autocorrelation Test - Ljung-Box test

The table below compiles the results of Ljung-Box tests carried out on different types of financial data: returns, squared returns and absolute returns, analysed at various time frequencies, i.e. daily, monthly and annual. These tests aim to identify the existence of autocorrelation in the data, a phenomenon where values in a time series are correlated with their own past values at different intervals or 'lags'. These results provide a basis for understanding the dynamics of time dependence and volatility of returns, which are essential for financial modelling and risk management.

Table 6: Combined Results of Ljung-Box Tests

	Lag 1	**Lag 2**	**Lag 3**
Returns			
Daily	6.14 (p=0.01)	6.15 (p=0.05)	6.98 (p=0.07)
Monthly	0.59 (p=0.44)	0.60 (p=0.74)	2.78 (p=0.43)
Annually	1.59 (p=0.21)	7.49 (p=0.02)	10.04 (p=0.02)
Squared Returns			
Daily	277.47 (p=0.00)	654.93 (p=0.00)	858.92 (p=0.00)
Monthly	0.00 (p=1.00)	0.76 (p=0.68)	0.77 (p=0.86)
Annually	0.06 (p=0.81)	0.24 (p=0.89)	0.71 (p=0.87)
Absolute Returns			
Daily	222.14 (p=0.00)	433.34 (p=0.00)	624.36 (p=0.00)
Monthly	0.02 (p=0.88)	1.02 (p=0.60)	1.07 (p=0.78)
Annually	0.20 (p=0.66)	0.28 (p=0.87)	1.20 (p=0.75)

3.6.1 Ljung-Box test formula

The test formula is as follows:

$$Q = n(n+2)\sum_{k=1}^{h}\frac{\hat{\rho}_k^2}{n-k}$$

where:

- n represents the number of observations in the time series.
- h is the number of lags for which autocorrelations are calculated.
- $\hat{\rho}_k$ is the autocorrelation estimated at lag k.

3.6.2 Interpretation of results

The daily results of the Ljung-Box test for returns, squared returns and absolute returns reveal significant insights into the autocorrelation and volatility of the data analysed.

For daily returns, the Ljung-Box test values indicate significant autocorrelation at lags 1 and 2 with p-values of 0.01 and 0.05, respectively. This suggests that one day's returns have some dependence on the returns of previous days. However, at lag 3, this autocorrelation becomes insignificant (p-value of 0.07), which could indicate that the influence of previous days' returns diminishes over time.

The results for daily squared returns show extremely high test statistics with p-values of zero at all three lags. This confirms the presence of volatility clusters in the data. In other words, days with large variation tend to be followed by other days with large variation, indicating significant persistence in the volatility of returns. This observation is crucial for financial modelling, as it suggests that a model capable of capturing this volatility, such as a GARCH model, would be appropriate.

Finally, daily absolute returns also show significant autocorrelation with p-values of zero for the three lags tested. This result confirms that the magnitude of variations in returns also exhibits significant autocorrelation, reinforcing the idea that periods of high volatility are clustered.

In conclusion, these analyses indicate that not only returns, but also their volatility and the magnitude of their variations are autocorrelated at the daily level. This justifies the use of models such as GARCHs to better forecast future returns and their volatility, taking into account both auto-correlation in returns and persistence in volatility variations. These results also highlight the importance of models that can adapt to conditional volatility for more effective risk management.

3.6.3 Cluster of volatility

When analysing the graphs of Novartis log returns and volatility, a visual observation of the data can enrich the statistical interpretations provided by tests such as the Ljung-Box.

By observing the graphs, it is possible to see clusters of volatility. These periods of high volatility, such as those seen around 2008 and 2020, coincide with major financial crises or significant events affecting the company, which could logically lead to autocorrelation in daily returns.

The six-month moving average chart clearly illustrates these periods of increased volatility and suggests that returns are not independent from one day to the next, but rather that volatility in one period tends to influence that of subsequent periods. This visual observation of the data therefore supports the idea that the results of the Ljung-Box test indicating autocorrelation are not simply a statistical artefact, but reflect real market behaviour.

Thus, by combining the visual analysis with the results of the Ljung-Box test, we can visually confirm the existence of autocorrelation in Novartis returns.

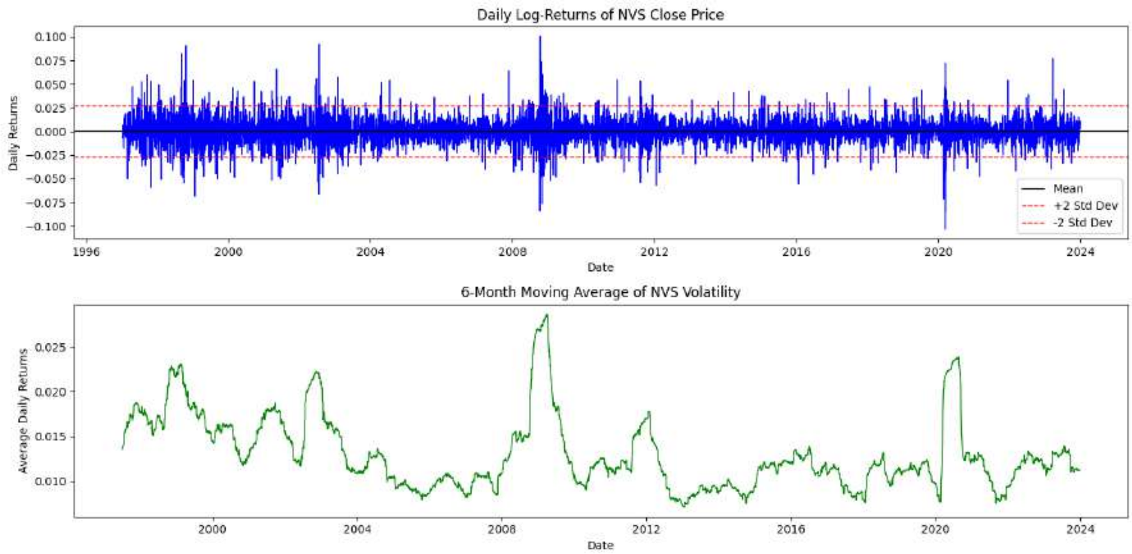

Figure 13: NVS Volatility

3.7 GARCH Model settings

A GARCH model, which stands for generalised conditionally heteroscedastic autoregressive model, is a statistical tool used mainly to analyse and predict the volatility of financial time series, such as share prices or exchange rates. This model is an extension of the ARCH model (conditionally heteroscedastic autoregressive model) developed by Robert Engle, for which he received the Nobel Prize in Economics. The GARCH model allows the variance of a time series to be modelled as a function of its past values and past variances.

In the context of a GARCH(p, q) model, the parameters 'p' and 'q' play crucial roles. The parameter 'p' represents the number of autoregressive variance terms, i.e. it indicates how many previous periods of the variance series are used to predict the current variance. On the other hand, the 'q' parameter refers to the number of moving average terms, including how many previous periods of squared errors (which reflect shocks or innovations in the variance) are taken into account in the model. The combination of these two parameters enables the GARCH model to capture both the persistence in volatility and the effects of recent shocks, providing a rich and flexible representation of the temporal dynamics of volatility.

The formula for a GARCH(p, q) model is given by :

$$\sigma_t^2 = \omega + \sum_{i=1}^{p} \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2 \quad (1)$$

where :

- σ_t^2 is the variance conditional at time t, representing the expected variance based on the information available up to $t-1$.
- ω is a constant term which ensures that the conditional variance is always positive.
- $\sum_{i=1}^{p} \alpha_i \epsilon_{t-i}^2$ is the sum of the last p squared error terms, weighted by the α_i parameters. These terms reflect the effect of past shocks on the current variance.
- $\sum_{j=1}^{q} \beta_j \sigma_{t-j}^2$ represents the sum of the last q conditional variances, weighted by the β_j parameters, making it possible to measure the persistent effect of past variances on the current variance.

The use of the Ljung-Box test played a crucial role in our analysis of Novartis returns, highlighting the need to consider autocorrelation, in particular volatility clustering, when modelling these data.

Given these results, it becomes imperative to model this autocorrelation in order to obtain more accurate and efficient forecasts. Consequently, we plan to adopt an approach based on several GARCH models. Each GARCH model will be configured with different orders of autoregressive terms and moving averages, allowing us to compare their ability to model the volatility of returns.

To determine the most appropriate model, we will use the Akaike Information Criterion (AIC). The AIC measures both the goodness of fit of the model and its complexity, penalising models that use an excessive number of parameters. A lower AIC indicates a better balance between the complexity of the model and its fit to the observed data. By selecting the GARCH model with the lowest AIC, we will be able to optimise our modelling of Novartis returns while taking into account the autocorrelation demonstrated by the Ljung-Box test.

Table 7: AIC values for different GARCH models as a function of orders p and q

AIC of GARCH(p,q)	p=1	p=2	p=3
q=1	-40716.35	-40693.80	-40688.68
q=2	-40703.78	-40691.75	-40678.95
q=3	-40748.63	-40725.74	-40708.94

The table presented shows the AIC values for different GARCH models with varying combinations of p (order of volatility autoregressive terms) and q (order of innovation moving averages).

Looking at the data, the GARCH(1,3) model has the lowest AIC (-40748.63), closely followed by the GARCH(2,3) model with an AIC of -40725.74. Although the GARCH(2,3) model shows a good fit, its AIC is slightly higher, indicating potential complexity without proportional improvement in fit compared to the GARCH(1,3).

Thus, based on the AIC criterion and the values provided, the GARCH(1,3) model should be chosen to model Novartis returns, as it offers the best balance between data fit and model complexity.

This graph shows a comparison between the observed daily volatility of Novartis returns (in pink) and the conditional volatility estimated by the GARCH(1,3) model (in blue) from 1997 to 2024. The GARCH(1,3) model is applied to historical data to estimate the conditional volatility of share returns.

Observation of the graph reveals that the GARCH(1,3) model succeeds in following the general trend in daily volatility, capturing periods of lower volatility particularly well. This ability to model volatility in calm periods is crucial for effective risk management in finance. The model also shows some responsiveness to volatility spikes, although it does not always capture the maximum amplitude of

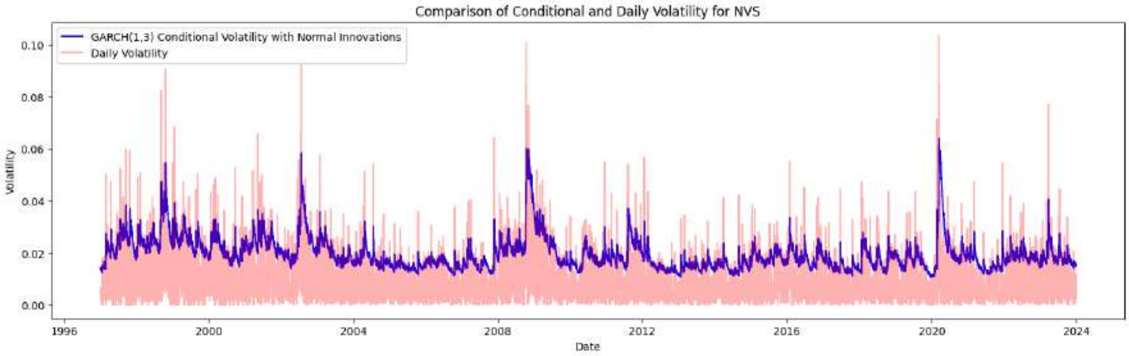

Figure 14: Conditionnal Volatility

these spikes, which is typical of GARCH models due to their conditional nature and dependence on past information.

In addition, areas where daily volatility greatly exceeds the modelled conditional volatility could indicate external events or market shocks not anticipated by the model based solely on historical data. This illustrates an inherent limitation of GARCH models: their ability to forecast is strongly conditioned by the volatility patterns already observed and less effective in situations of market disruption or significant macroeconomic announcements.

In sum, the GARCH(1,3) model demonstrates a good ability to track the dynamics of Novartis volatility over the long term, providing a valuable tool for risk assessment and financial decision-making. However, ongoing monitoring and adjustments to the model may be required to maintain its relevance in the face of changing market conditions.

3.8 Monte-Carlo simulation on Returns

As part of our study, we have implemented a Monte Carlo simulation comprising 10,000 trajectories to analyse the evolution of a share price. The starting point of this simulation is the share price at the close of 2023, and it extends over 252 trading days, roughly the equivalent of one trading year (to reach the close of 2024).

To strengthen the reliability of our projections, various statistical parameters have been carefully adjusted. In particular, we selected the GARCH(1,3) model after testing several variants for modelling return volatility. In addition, we parameterised a skewed Student's t distribution to better reflect the asymmetries and thick tails observed in real return distributions.

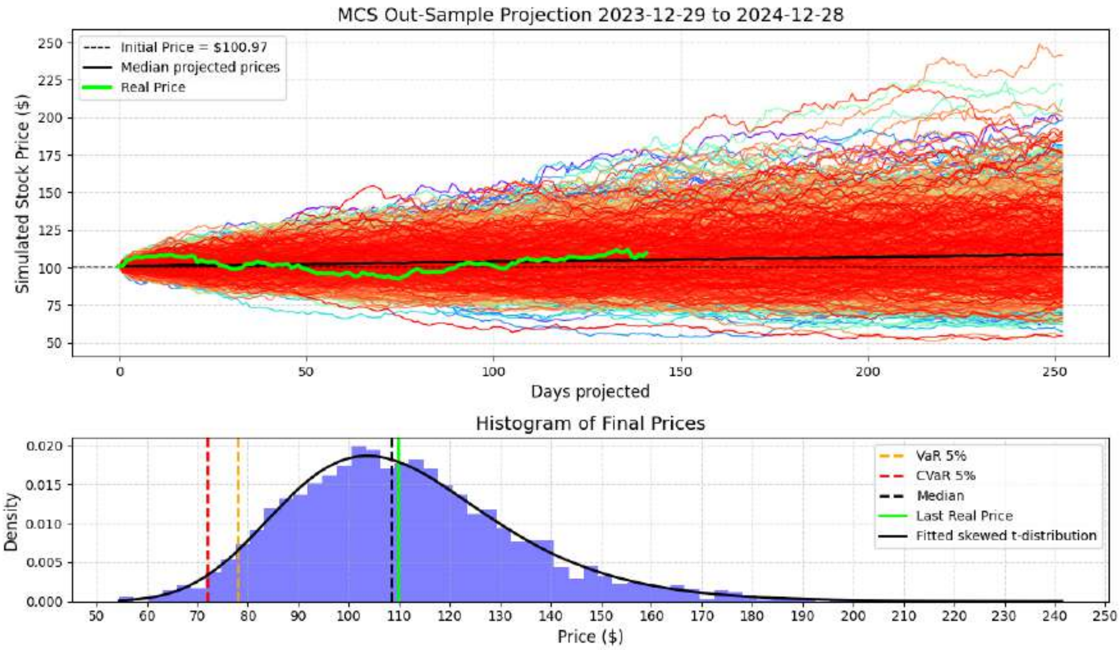

Figure 15: MCS on NVS

On the simulation chart, the median of projected prices is drawn as a solid black line. This median line helps us to understand the central tendency of the projections over time. At the same time, the actual Novartis share price is indicated by the green line, providing a direct point of comparison with our simulations.

Below the main graph, a histogram of the final prices estimated by the Monte-Carlo simulation is presented. This histogram is fitted by an asymmetric Student's t distribution, which illustrates the distribution of possible outcomes at the end of the simulated period. This adjustment is essential for parametric risk management applications, enabling extreme risks and potential loss scenarios to be assessed more accurately.

Returns r_t follow a t-skewed distribution, where volatility is modelled by a GARCH model:

$$r_t \sim ST(\mu, \sigma(t), \gamma, \nu)$$

The relationship between the current price and the previous price using returns is given by :

$$P_t = P_{t-1} \times e^{r_t}$$

where P_t is the price at time t, P_{t-1} is the price at time $t-1$, and r_t is the log return.

3.8.1 Observations in relation to empirical averages and medians

In this Monte Carlo Simulation (MCS) analysis for the projection of a share price over a one-year period, the median of the projected trajectories follows an apparently linear trend. This median projection does not coincide with the values obtained by conventional statistics, such as empirical annualised returns, whose averages and medians are 5.091% and (0.025*252 =) 6.3%, respectively. It also diverges from the results obtained by calculations based on an annual frequency, with values of 3.918%. and 5.872%.

This study highlights the importance of numerical simulations in assessing equity returns. The median of the simulated prices, reflecting a return of 8.46%, demonstrates that numerical simulations can reveal price behaviour that is not evident via traditional statistical analysis. These simulations incorporate a wide range of possible scenarios, offering a more comprehensive estimate of potential risks and returns, which diverge significantly from approaches limited to historical data.

The use of these simulations is essential for investors and analysts, as they provide an in-depth perspective on potential volatility and price movements, essential for risk management and strategic decision-making. The graph of the histogram of simulated final prices and the distribution fitted to an asymmetric Student's t distribution provides an effective way of visualising the dispersion of simulated results. This visualisation highlights the risks and asymmetric returns that simple mean or median calculations fail to capture.

4 Statistics Final Price Distribution

The table below shows the statistics applied to the one-year Novartis share prices of the 2025 projections, obtained from the two Monte Carlo simulations. The results highlight the difference in risk.

Table 8: Statistics of final prices

Statistic	**MCS Stock Returns**	**MCS FCFF**
Mean	110.22	111.04
Median	109.51	106.62
Standard Deviation	22.03	32.44
Minimum	54.96	37.49
Maximum	199.34	335.30
IQR	28.11	41.18
Skewness	0.64	0.94
Kurtosis	1.13	1.81
VaR 95%	77.74	66.33
CVaR 95%	70.72	58.83

The difference between the mean and the median in the MCS results for FCFF is more pronounced than in the MCS results for stock returns. This difference is a proxy for skewness. Other statistics, such as standard deviation, maximum and minimum values, and risk metrics (VaR and CVaR), show that the model based on FCFF includes greater variability and uncertainty. The significantly higher standard deviation in the FCFF model indicates a greater dispersion of results around the mean, which means that forecasts can vary considerably from one scenario to another. In addition, the lower VaR and CVaR show that in the worst-case scenarios, the potential losses are greater, underlining a greater view of risk.

It is important to note that the FCFF model is less robust than the returns model, largely because it was calibrated with a limited amount of available data. In other words, the model has been heavily influenced by extreme values in FCFF. This results in modelling of fairly strong variations in FCFF, which can lead to less reliable forecasts. This reduced robustness is problematic because it can make the model more vulnerable to error or bias.

5 Backtesting In-Out Sample

5.1 Methodology

In this study, the In-Out Sample backtesting methodology is used to assess the accuracy of the statistical models used, by adapting the sample period to match the data available up to one year before our estimation point. The actual Novartis price used is that observed at the projection date, i.e. one year after the end of the last sample date (backtesting covers the period from 2011 to 2024) This process focuses on the comparison between the median of the stochastic projections and the actual Novartis share price, as well as the average of the prices projected by nine major investment banks (Barclays, Bank of America, Deutsche Bank, Jefferies, JP Morgan Chase, Morgan Stanley, Société Générale, UBS, and Vontobel). These banks were chosen solely on the basis of the data available on Refinitiv Eikon.

The aim is to compare these predictions with actual market and analyst performance and to provide a comparative perspective on the accuracy of different estimation methods. All the backtesting results tables are available in the appendices, providing a detailed analysis of comparative performance over the test period from 2010 to 2024.

5.2 Results

The results of the study indicate that the two models used underestimate and overestimate the real price in 50% of cases respectively, suggesting that the median of the estimates is correctly placed. In contrast, analysts tend to overestimate the price in 73% of situations, revealing a more pronounced optimism in their projections.

To observe the accuracy of the three estimation methods, we will use the average absolute deviations between the estimated price and the actual price. The use of average absolute deviations between the estimated price and the actual price is essential because this measure eliminates the risk of compensation between negative and positive errors. This is because, unlike deviations which can cancel each other out when averaged, absolute deviations maintain the integrity of each error, thus faithfully reflecting the total magnitude of the deviations without directional bias.

These deviations are as follows: 14.6% for the analysts' target prices, 11.4% for the median of the share price projections, and 10.4% for the estimate based on the median of the FCFF projections.

The relative stability of the Novartis share price favours these models, particularly in contexts where projections tend to underestimate the performance of companies experiencing recent and marked growth.

When the share price for the year is stable with a slight increase, the models, based on average upward trends, anticipate the price fairly well, approaching the median or mean of the distribution. It appears that the average projections are more conservative estimates than those of analysts.

The statistical analysis underlines that if the model is well parameterised and shows a strongly leptokurtic distribution with low assymetry, estimation by the median or mean as a 'point' value of estimation would make sense. However, there is too much uncertainty, even on a company considered as stable, to adopt this point of view.

There is a difference between the median prices of the two models, with the FCFF model systematically showing a slightly lower median price than the MSC model, which could be linked to the assumptions concerning the terminal growth rate.

It would be advisable to carry out a sensitivity analysis on this terminal growth rate to better understand its impact on the estimates, thus requiring a three-dimensional visualisation to illustrate these dynamics.

5.3 Meaning on the market efficiency hypothesis

We find that the use of a statistical model, fed solely by two internal company time series (stock price and FCFF), based on empirical finance approaches, offers neither better nor worse significant performance in estimating the future share price (calculated from the median of the projected stochastic paths) compared with the average target price established by analysts. The results of this study seem to corroborate the semi-strong efficient market hypothesis.

This hypothesis holds that financial asset prices tend to reflect quickly and completely all available public information, such as market data, financial reports and news. Empirical studies tend to corroborate this idea, suggesting that it is difficult for an investor to outperform the market by relying solely on this information, as it is generally already incorporated into prices. However, there are price anomalies, particularly during periods of speculation or bubble formation, which show that markets can sometimes deviate from this assumption, offering atypical opportunities for gains.

5.4 Complementarity with analysts

Despite the fact that stochastic projections can sometimes match or approach the analysts' estimates, the role of analysts remains essential in the financial ecosystem for several reasons:

1. **Synthesis of complex information** : Analysts reduce a wide range of financial, economic and sectoral information into accessible recommendations, helping investors to make informed and reasoned decisions.

2. **Sectoral expertise** : They offer valuable insights based on in-depth knowledge of the sectors and companies they follow, a dimension that stochastic models do not always fully capture.

3. **Long-term outlook** : Analysts often take into account long-term prospects and corporate strategies, not just short-term price movements.

4. **Market Influence** : Analysts' opinions can significantly influence market perceptions and stock valuations, playing an active role in market dynamics.

5. **Diversity of opinion** : Analysts' estimates provide a diversity of viewpoints which enrich the debate and investment decisions on the financial markets.

In short, stochastic projections, while valuable for their statistical approach, complement rather than replace analysts' assessments. Analysts bring a human and contextual perspective that remains invaluable for investment decisions, illustrating the importance of combining quantitative and qualitative analyses for a complete view of the market.

5.5 Analyst bias

However, it is important to bear in mind the presence of 'analyst bias', a systematic tendency for financial analysts to make potentially biased predictions or recommendations. This bias can be influenced by a variety of factors and can sometimes lead to less objective or inaccurate assessments of the market. Here are some common sources of analyst bias:

1. **Confirmation bias**: Analysts may favour information that confirms their pre-existing beliefs or previous analyses, downplaying or ignoring data that could contradict their views.

2. **Conflicts of interest**: Analysts working for institutions that have commercial relationships with the companies they cover may be influenced to give more favourable recommendations. For example, if a bank hopes to win business from a company, its analysts may be encouraged to give a positive assessment of that company's shares.

3. **Optimism bias**: Research shows that analysts are often over-optimistic in their forecasts of earnings and share performance. This may be due to a relationship with the management of the companies analysed or a tendency to maintain good relations with these companies.

4. **Institutional pressures**: Analysts may be under pressure from their employers to produce reports that support the institution's business objectives or trading strategies.

5. **Recency bias**: This occurs when analysts give too much weight to recent events at the expense of longer-term trends or broader contexts, which can lead to recommendations that are not aligned with a company's long-term prospects.

6. **Herd effect**: Some analysts may follow majority opinions or the recommendations of other analysts without conducting a full independent assessment, to avoid deviating significantly from the consensus, which may be perceived as an occupational hazard.

In our analysis, we identified trends corresponding to these biases, in particular excessive optimism and too strong a consensus, manifested by a normal distribution with a surprisingly low standard deviation.

6 Conclusion

In this study, we have attempted to establish the distribution of future Novartis prices for the end of 2024. To do this, we empirically analysed available data from 1997 to 2023. Two approaches were adopted: the first involved projecting future FCFF, discounting them using the DCF method and deriving final prices. The second approach involved modelling stock market returns to derive final prices. Both methods used Monte Carlo simulations.

This work illustrated that the actual daily stock return, denoted r_t, corresponds to a random variable whose distribution law, the number of parameters and the values of these parameters remain undetermined. Moreover, the complexity is increased by the fact that these parameters vary daily.

On the basis of projections of these daily returns, it is possible to determine the final prices one year ahead, thus arriving at a distribution of final prices for this period. Our estimates, based on numerical simulations, suggest that this distribution may approach a skewed-t.

6.1 Distribution of actual daily returns and one-year price

The final one-year price of the Novartis share (P_{252}) is a random variable that is mathematically determined by the daily stock market returns of Novartis, which are also random variables whose parameters are dynamic over time. This can be formulated as follows:

$$P_{252} = P_0 \times \prod_{t=1}^{252} e^{r_t}, \quad r_t \sim D(\theta_t)$$

- where θ_t is a vector of unknown time-dependent parameters,
- and D is an unknown distribution.

The actual distribution of the final price over a year, which is impossible to determine exactly, can be written as :

$$P_{252} \sim D(\Theta)$$

where Θ is a distinct parameter vector, different from the θ_t used for daily returns.

6.2 Our approximation of real distributions

We have modelled the daily returns of Novartis according to a skewed-t distribution with dynamic volatility using a GARCH model. This is presented as follows:

The return r_t follows a Skewed-t distribution with the following parameters:

$$r_t \sim \text{Skewed-t}(\mu, \sigma_t, nc, df)$$

where

- μ is the mean of the distribution,
- σ_t is the standard deviation, which may vary over time,
- nc is the asymmetry coefficient,
- df is the degree of freedom of the distribution

The final one-year price, P_{252} was approximated in both methods via a Skewed-t distribution according to the following equation:

$$P_{252} \sim \text{Skewed-t}(\mu, \sigma, nc, df)$$

6.3 Distribution of analysts' target prices

In the field of financial analysis, analysts often establish 'target prices' for shares, representing their estimates of the price of a share one year ahead. These estimates are based on various economic models and assumptions, and vary from one analyst to another depending on their interpretation of the data and the company's future prospects.

If we compile all these point estimates of one-year target prices provided by different analysts for a given share, it is possible to analyse them statistically to observe the distribution of these estimates. Theoretically, if the number of estimates is large enough and the individual biases of the analysts are independent of each other, these target prices could tend towards a normal distribution. This normality would be due to the centralisation of the different valuations around a mean, with a standard deviation that reflects the dispersion of opinions.

It is assumed that 252-day target prices follow a normal distribution, with the following parameters:

$$\text{'Target prices'} \sim \mathcal{N}(\mu, \sigma^2)$$

where

- μ is the average of the target prices,
- σ^2 is the variance of these prices.

6.4 Comparison of one-year price distributions

Given that it is intrinsically impossible to predict the exact price of a share one year from now, because fluctuations in the financial markets, influenced by a multitude of economic, political and social factors, make each future price the outcome of a random process, attempting to estimate a specific price at a future date is not only imprecise but also uninformative. This is often reflected in the practice of analysts formulating 'target prices', which are supposed to represent future estimates. Although common, this approach does not capture the complexity and uncertainty inherent in the markets.

In the backtesting for this study, we in turn considered establishing a one-year target price for Novartis shares using the median of projected prices. However, we realised that these approaches, while popular, lack a robust scientific basis in volatile and unpredictable markets. Rather than providing a precise price, we need to move towards a final price distribution approach.

We have focused on a more robust and informative approach by modelling the one-year price distribution. This method does not allow us to predict an exact value, but to understand and characterise all possible outcomes, encapsulated in a probability distribution. Using this model, we can estimate the probability that a share price will fall within a certain range, offering a richer, more nuanced perspective on future market expectations.

The projected price, as a random variable, is best understood through the parameters of its probability distribution, such as mean, variance, skewness and kurtosis. These parameters not only provide information on the expected average price, but also on volatility and the risk of extreme prices. Ultimately, modelling this distribution allows investors and analysts to better manage risk and develop more sophisticated investment strategies tailored to the risk profile.

This research therefore confirms that the best way to understand and forecast future share prices is not in the quest for precision, but in the precise definition and in-depth understanding of the probability distribution of prices, which represents as closely as possible the true random and uncertain nature of financial markets.

6.4.1 Impossibility of determining the quality of the models

From a formal and theoretical perspective, our objective was to maximize the area of intersection under the curves between our estimated model of the annual distribution, represented by a Skewed-T distribution $ST(x;\theta)$, and the true distribution $D(x;\Theta)$.

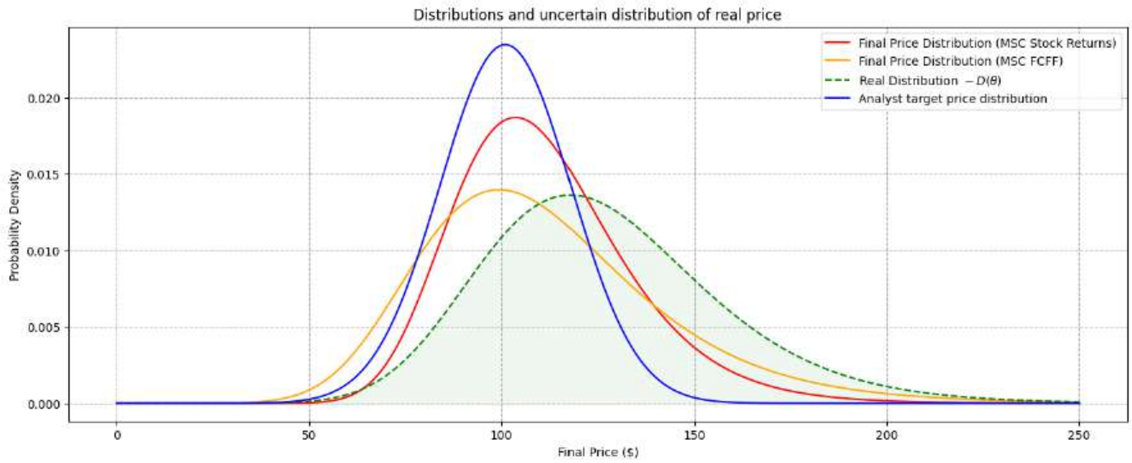

Figure 16: All pdf of final prices

Maximizing the intersection area of two parametric distributions is mathematically formulated as follows:

$$\text{Maximize} \quad \int_{-\infty}^{\infty} \min(ST(x;\theta), D(x;\Theta))\, dx$$

However, given the lack of exact knowledge of the true distribution, our approach effectively translates into an attempt to optimize the plausibility of our estimate. In other words, we seek to adjust our model so that it best approximates what we hypothesize to be the underlying true distribution of the data.

6.5 Taking speculative positions on statistical convictions

A speculative position should be based on the mathematical expectation of gain, as proposed by Louis Bachelier in his 'Théorie de la Spéculation', and not on the median of stochastic projections or a target price derived from a single DCF model. According to this theory, the stock price should be determined by the mathematical expectation of its future price, i.e. the average of possible future prices weighted by their probabilities.

In practice, however, we do not have the exact distribution of the 'true price' in one year's time. In the absence of this precise distribution, it is necessary to rely on one's own assessment of the one-year price curve to estimate this mathematical expectation.

Thus, the decision to invest should be based on a comparison between the expected price in one year's time, according to the investor's own estimate ($E[P_{t+252}]$), and the current price (P_t). If the mathematical expectation of the future price, according to this estimate, is higher than the current price, the mathematical expectation is positive, which suggests a long (buy) position in the asset. On the other hand, if the mathematical expectation of the future price is lower than the current price, the mathematical expectation is negative, suggesting a short (sell) position.

Therefore, the mean, $E(X) = \int_{-\infty}^{\infty} x f(x)\, dx$, (and not the median) should be considered as an **indicator of mathematical advantage** (because it encapsulates all the risks by weighting them in a single value) and **in no way as a punctual valuation**.

References

[1] Eric Jondeau, Ser-Huang Poon, Michael Rockinger. *Financial Modeling Under Non-Gaussian Distribution*. Springer, 2007.

[2] Paolo Brandimarte. *Handbook in Monte Carlo Simulation - Applications in Financial Engineering, Risk Management, and Economics*. Wiley, 2014.

[3] Yves Hilpisch. *Python for Finance Analyze Big Financial Data*. O'REILLY, 2015.

[4] Jon Danielsson. *Financial Risk Forecasting*. Wiley, 2011.

[5] Louis Bachelier. *Théorie de la spéculation*. É.N.S., 1900.

[6] Davis W. Edwards. *Risk Management in Trading*. Wiley, 2014.

Appendix

Density Probability for the Skewed-t Distribution

The probability density for a non-central t distribution is given by :

$$f(x \mid df, \mu, \sigma) = \frac{df^{df/2}\Gamma(df+1)}{\sqrt{2\pi}\Gamma(df/2)\sigma^{df}} \left(1 + \frac{(x-\mu)^2}{df\sigma^2}\right)^{-(df+1)/2} \exp\left(-\frac{\mu^2}{2\sigma^2\left(1+\frac{(x-\mu)^2}{df\sigma^2}\right)}\right)$$

where :

- x is the random variable,
- df is the degree of freedom,
- μ is non-centrality,
- σ is the scale,
- Γ is the gamma function.

WACC Formula

The WACC is calculated by weighting the costs of equity and debt by their respective proportions in the company's financing structure:

$$WACC = \left(\frac{E}{V} \times R_e\right) + \left(\frac{D}{V} \times R_d\right)$$

where,

- E is the market value of the equity
- D is the market value of the debt
- V is the sum of E and D, so 100%.

Cost of Equity (CAPM Model)

The cost of equity (R_e) is calculated using the CAPM model, defined by the formula:

$$R_e = R_f + \beta \times (ERP) + CRP$$

where,

- R_f is the risk-free rate
- β is the leverage beta
- ERP is the market risk premium
- CRP is the country risk premium

Significance tests for GARCH models

Table 9: Significance of parameters for GARCH(1,1), GARCH(1,2), and GARCH(1,3)

Coefficient	Estimate	P-Value	Significant at 5%
	GARCH(1,1)		
ω	0.0000	0.0000	True
α_1	0.0498	0.0000	True
β_1	0.9263	0.0000	True
η	6.0859	0.0000	True
λ	-0.0040	0.7926	False
	GARCH(1,2)		
ω	0.0000	0.0000	True
α_1	0.0500	0.0000	True
β_1	0.4650	0.0277	True
β_2	0.4650	0.0263	True
η	6.0495	0.0000	True
λ	0.0000	1.0000	False
	GARCH(1,3)		
ω	0.0000	0.0000	True
α_1	0.0789	0.0000	True
β_1	0.4496	0.0030	True
β_2	0.0543	0.7510	False
β_3	0.4050	0.0071	True
η	5.0520	0.0000	True
λ	-0.0436	0.0045	True

Table 10: Significance of parameters for GARCH(2,1), GARCH(2,2), and GARCH(2,3)

Coefficient	Estimate	P-Value	Significant at 5%
	GARCH(2,1)		
ω	0.0000	0.0000	True
α_1	0.0250	0.0185	True
α_2	0.0250	0.0198	True
β_1	0.9300	0.0000	True
η	6.1031	0.0000	True
λ	0.0000	1.0000	False
	GARCH(2,2)		
ω	0.0000	0.0000	True
α_1	0.0250	0.0224	True
α_2	0.0250	0.0440	True
β_1	0.4650	0.2231	False
β_2	0.4650	0.2183	False
η	6.0436	0.0000	True
λ	0.0000	1.0000	False
	GARCH(2,3)		
ω	0.0000	0.0000	True
α_1	0.0494	0.0001	True
α_2	0.0503	0.0004	True
β_1	0.2899	0.0102	True
β_2	0.2899	0.0872	False
β_3	0.2974	0.0058	True
η	5.8297	0.0000	True
λ	-0.0096	0.5355	False

Table 11: Significance of parameters for GARCH(3,1), GARCH(3,2), and GARCH(3,3)

Coefficient	Estimate	P-Value	Significant at 5%
	GARCH(3,1)		
ω	0.0000	0.0000	True
α_1	0.0176	0.0639	False
α_2	0.0176	0.2939	False
α_3	0.0172	0.2021	False
β_1	0.9262	0.0000	True
η	6.0771	0.0000	True
λ	-0.0050	0.7468	False
	GARCH(3,2)		
ω	0.0000	0.0000	True
α_1	0.0167	0.1768	False
α_2	0.0167	0.3714	False
α_3	0.0167	0.2312	False
β_1	0.4650	0.3580	False
β_2	0.4650	0.3529	False
η	6.0374	0.0000	True
λ	0.0000	1.0000	False
	GARCH(3,3)		
ω	0.0000	0.0000	True
α_1	0.0333	0.0001	True
α_2	0.0333	0.0010	True
α_3	0.0333	0.0430	True
β_1	0.2933	0.1986	False
β_2	0.2933	0.5394	False
β_3	0.2933	0.3216	False
η	5.9026	0.0000	True
λ	0.0000	1.0000	False

Results of backtesing (Target Price 01.2011)

Actual Price NVS (01.2011)		53$

Analysts	Target Price	Difference
Barclays	63$	19%
Bank of America	65$	23%
Deutsche Bank	62$	17%
Jefferies	65$	23%
JP Morgan Chase	69$	30%
Morgan Stanley	69$	30%
Société Générale	57$	8%
UBS	61$	15%
Vontobel	58$	9%
Mean	63$	19%

MCS on stock return	51$	-4%
MCS on FCFF	50$	-6%

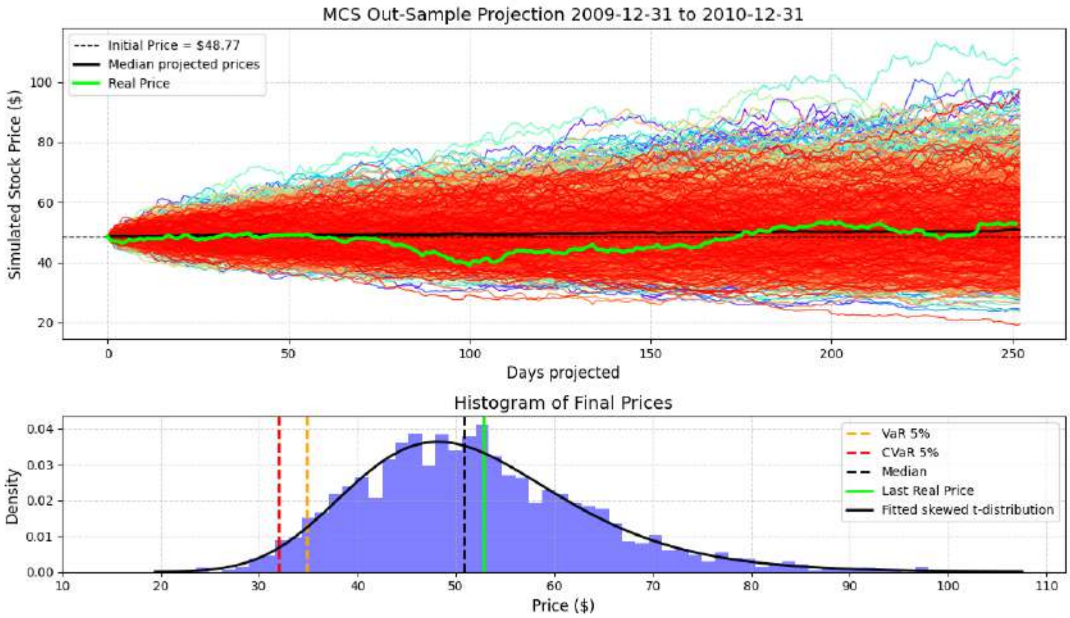

Results of backtesing (Target Price 01.2012)

Actual Price NVS (01.2012)	52$

Analysts	Target Price	Difference
Barclays	57$	10%
Bank of America	67$	28%
Deutsche Bank	60$	16%
Jefferies	64$	24%
JP Morgan Chase	61$	18%
Morgan Stanley	60$	16%
Société Générale	60$	16%
UBS	56$	8%
Vontobel	68$	31%
Mean	62$	19%

MCS on stock return	55$	6%
MCS on FCFF	53$	2%

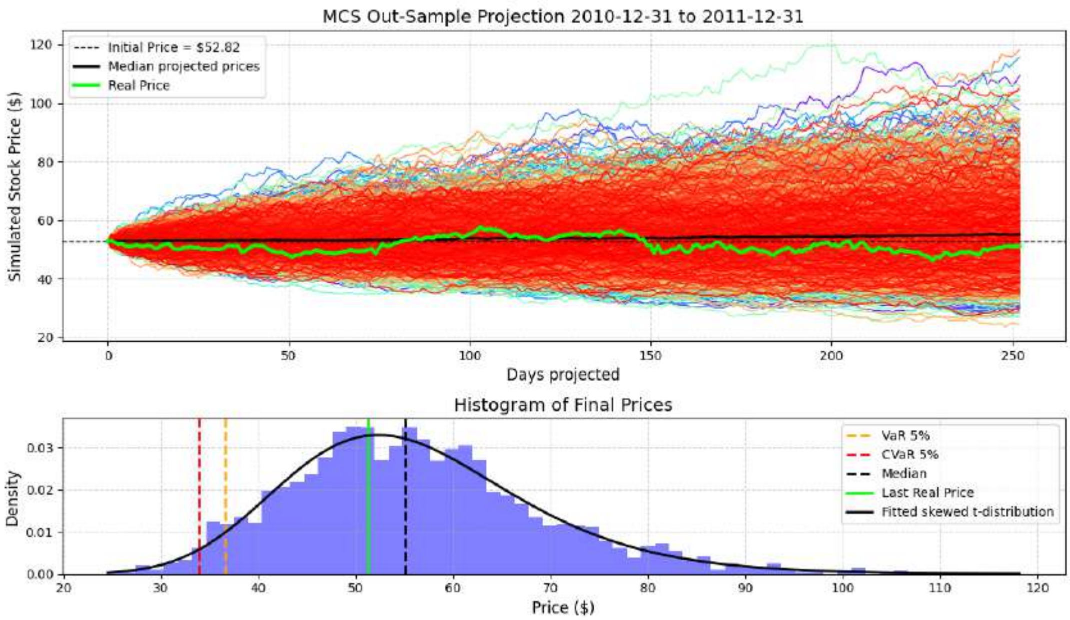

Results of backtesing (Target Price 01.2013)

Actual Price NVS (01.2013)	57$

Analysts	Target Price	Difference
Barclays	57$	1%
Bank of America	59$	3%
Deutsche Bank	60$	5%
Jefferies	69$	21%
JP Morgan Chase	63$	11%
Morgan Stanley	69$	20%
Société Générale	53$	-8%
UBS	50$	-13%
Vontobel	53$	-7%
Mean	59$	4%

MCS on stock return	54$	-5%
MCS on FCFF	53$	-7%

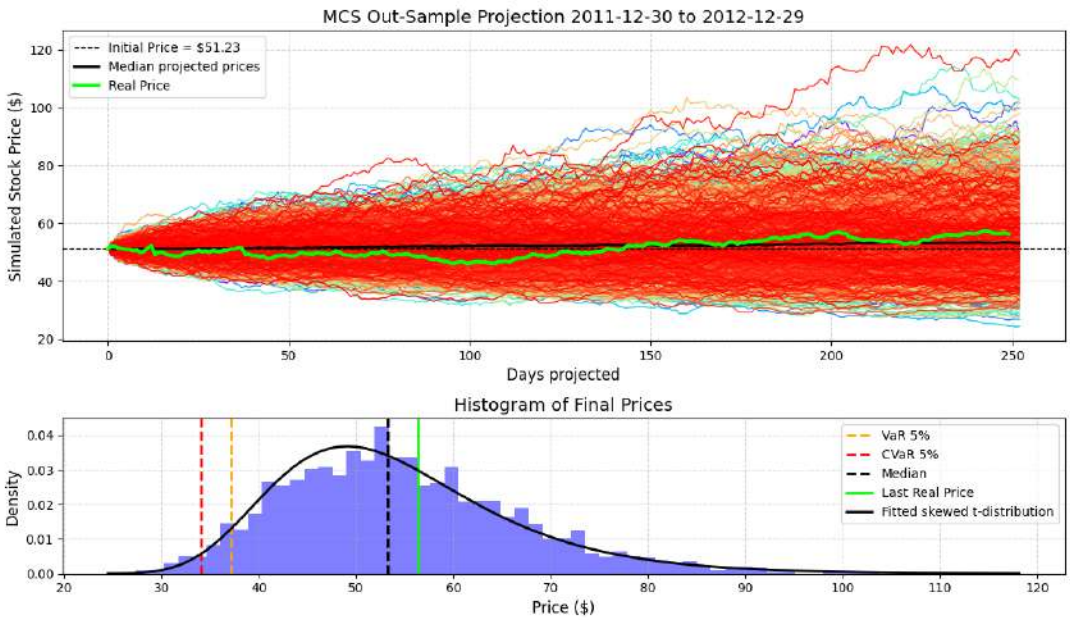

Results of backtesing (Target Price 01.2014)

Actual Price NVS (01.2014)	71$

Analysts	Target Price	Difference
Barclays	68$	-4%
Bank of America	60$	-15%
Deutsche Bank	72$	1%
Jefferies	62$	-13%
JP Morgan Chase	66$	-7%
Morgan Stanley	67$	-6%
Société Générale	71$	-1%
UBS	70$	-2%
Vontobel	69$	-3%
Mean	67$	-6%

MCS on stock return	60$	-15%
MCS on FCFF	58$	-18%

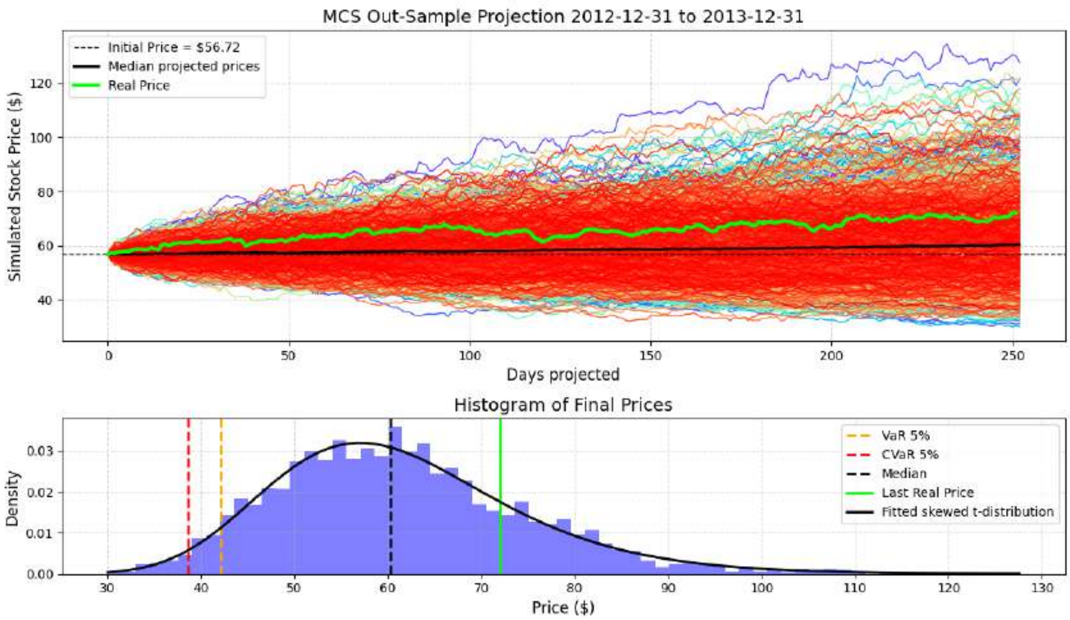

Results of backtesing (Target Price 01.2015)

Actual Price NVS (01.2015)	83$

Analysts	Target Price	Difference
Barclays	72$	-14%
Bank of America	81$	-2%
Deutsche Bank	75$	-10%
Jefferies	77$	-7%
JP Morgan Chase	75$	-10%
Morgan Stanley	74$	-11%
Société Générale	77$	-7%
UBS	80$	-3%
Vontobel	73$	-12%
Mean	76$	-8%

MCS on stock return	77$	-7%
MCS on FCFF	74$	-11%

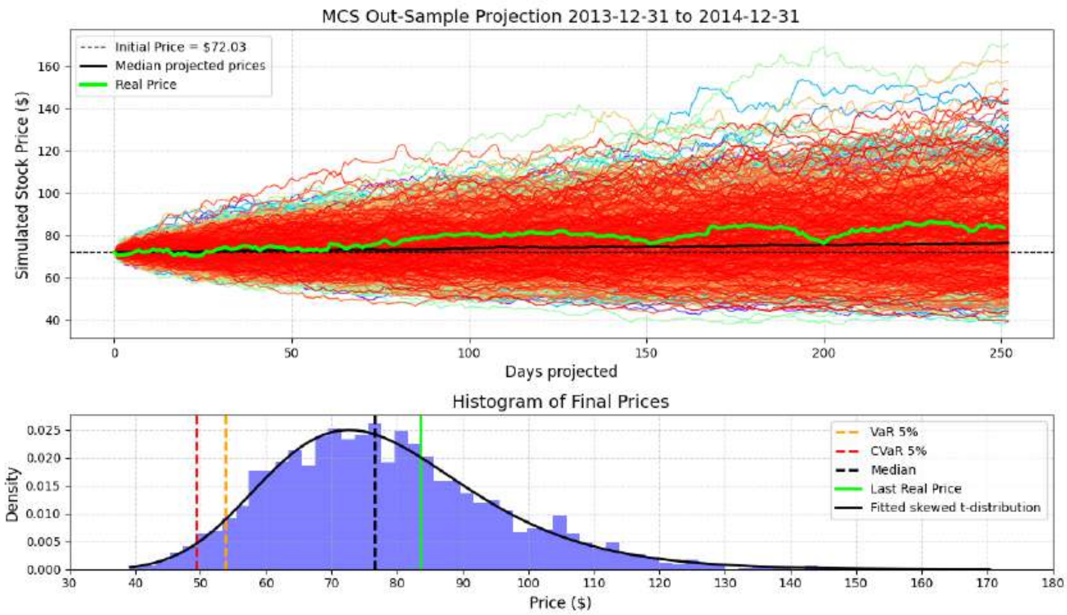

Results of backtesing (Target Price 01.2016)

Actual Price NVS (01.2016)	77$

Analysts	Target Price	Difference
Barclays	99$	29%
Bank of America	99$	29%
Deutsche Bank	100$	30%
Jefferies	85$	10%
JP Morgan Chase	117$	52%
Morgan Stanley	79$	3%
Société Générale	110$	43%
UBS	99$	29%
Vontobel	96$	25%
Mean	98$	27%

MCS on stock return	90$	17%
MCS on FCFF	85$	10%

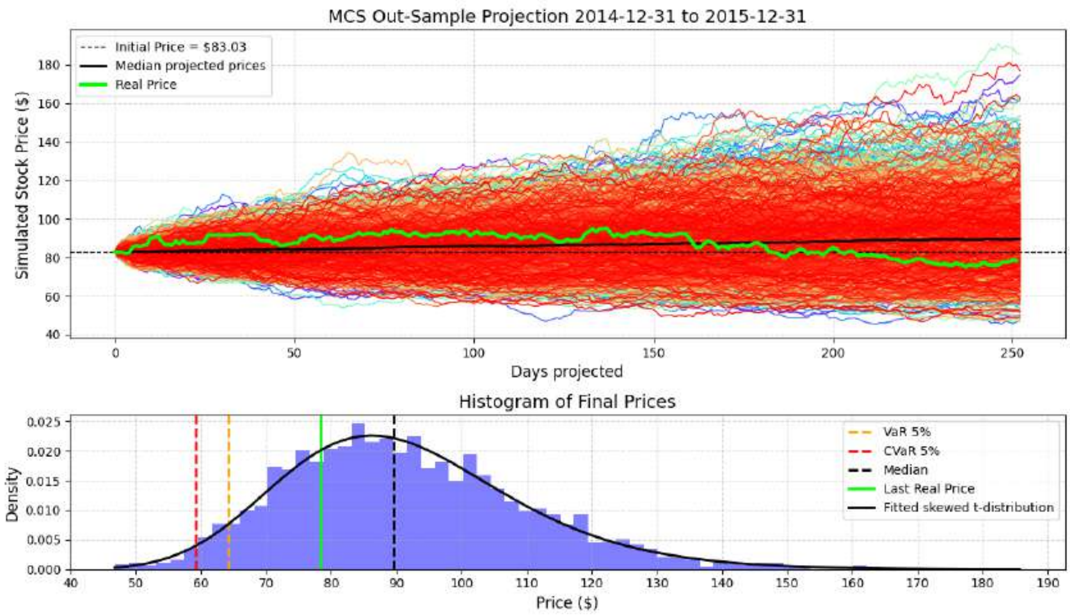

Results of backtesing (Target Price 01.2017)

Actual Price NVS (01.2017)	65$

Analysts	Target Price	Difference
Barclays	84$	29%
Bank of America	77$	19%
Deutsche Bank	98$	51%
Jefferies	95$	46%
JP Morgan Chase	84$	30%
Morgan Stanley	95$	46%
Société Générale	96$	48%
UBS	91$	41%
Vontobel	81$	24%
Mean	89$	37%

MCS on stock return	82$	26%
MCS on FCFF	80$	23%

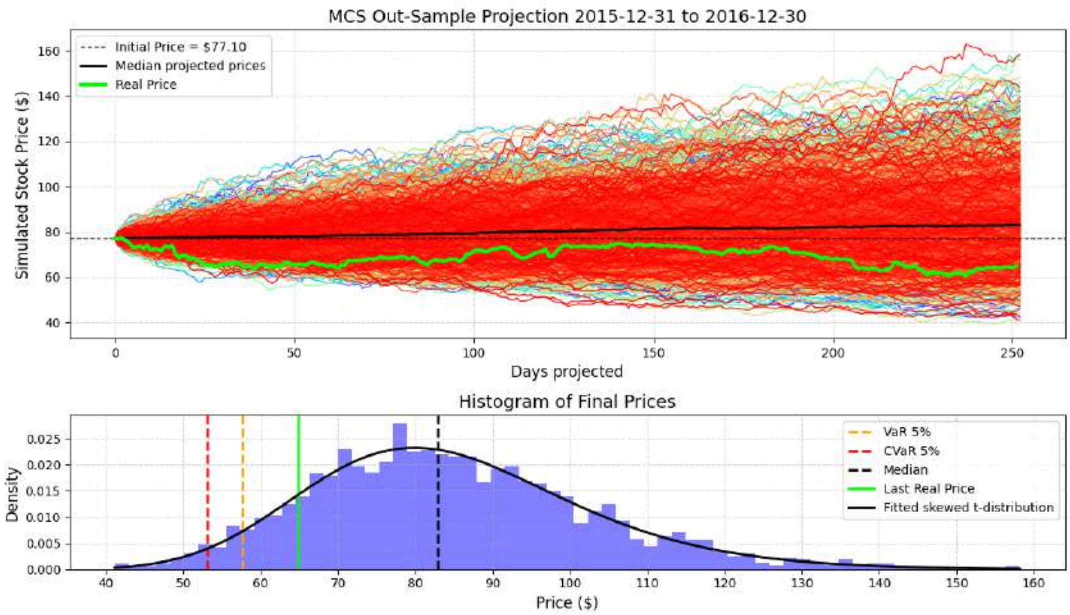

Results of backtesing (Target Price 01.2018)

Actual Price NVS (01.2018)		76$

Analysts	Target Price	Difference
Barclays	62$	-18%
Bank of America	65$	-15%
Deutsche Bank	95$	24%
Jefferies	66$	-13%
JP Morgan Chase	68$	-10%
Morgan Stanley	62$	-19%
Société Générale	50$	-35%
UBS	63$	-17%
Vontobel	74$	-3%
Mean	67$	-12%

MCS on stock return	69$	-9%
MCS on FCFF	67$	-12%

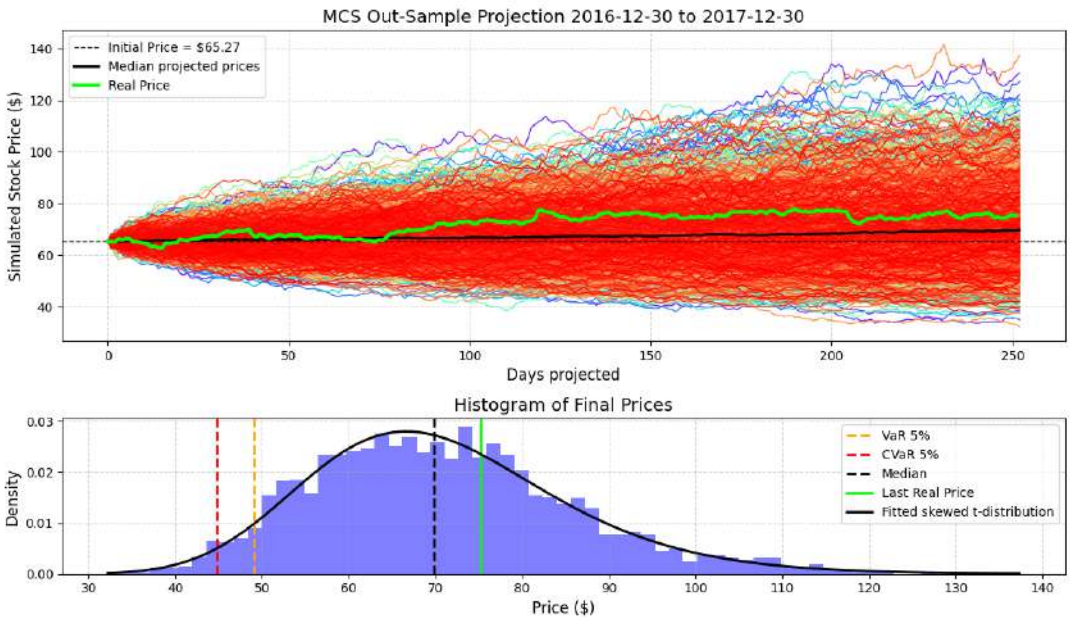

Results of backtesing (Target Price 01.2019)

Actual Price NVS (01.2019)	76$

Analysts	Target Price	Difference
Barclays	99$	30%
Bank of America	84$	11%
Deutsche Bank	67$	-12%
Jefferies	74$	-3%
JP Morgan Chase	88$	16%
Morgan Stanley	97$	28%
Société Générale	90$	18%
UBS	78$	3%
Vontobel	90$	18%
Mean	85$	12%

MCS on stock return	80$	5%
MCS on FCFF	77$	1%

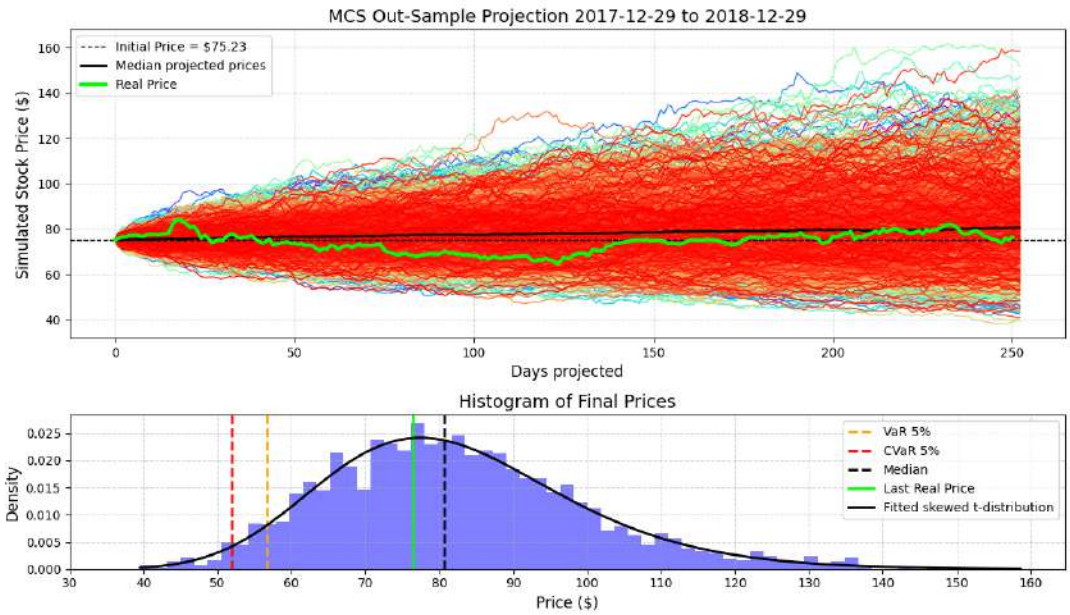

Results of backtesing (Target Price 01.2020)

Actual Price NVS (01.2020)	95$

Analysts	Target Price	Difference
Barclays	99$	4%
Bank of America	62$	-35%
Deutsche Bank	71$	-25%
Jefferies	87$	-8%
JP Morgan Chase	90$	-5%
Morgan Stanley	50$	-47%
Société Générale	74$	-22%
UBS	74$	-22%
Vontobel	75$	-21%
Mean	76$	-20%

MCS on stock return	82$	-14%
MCS on FCFF	80$	-16%

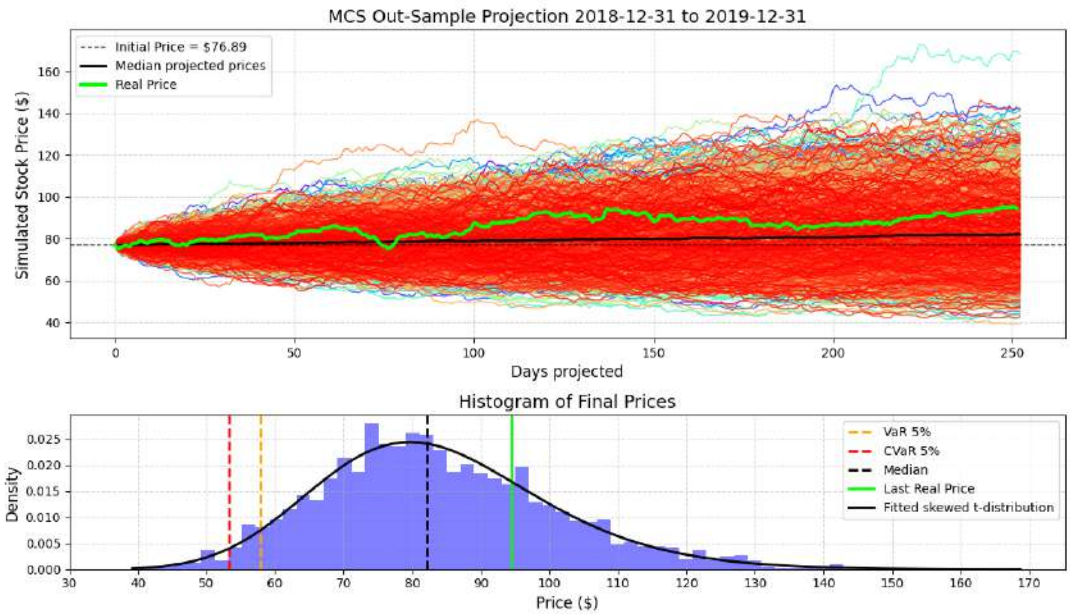

Results of backtesing (Target Price 01.2021)

Actual Price NVS (01.2021)	93$

Analysts	Target Price	Difference
Barclays	90$	-3%
Bank of America	85$	-9%
Deutsche Bank	95$	2%
Jefferies	99$	6%
JP Morgan Chase	99$	6%
Morgan Stanley	79$	-15%
Société Générale	92$	1%
UBS	101$	9%
Vontobel	109$	17%
Mean	94$	1%

MCS on stock return	102$	10%
MCS on FCFF	98$	5%

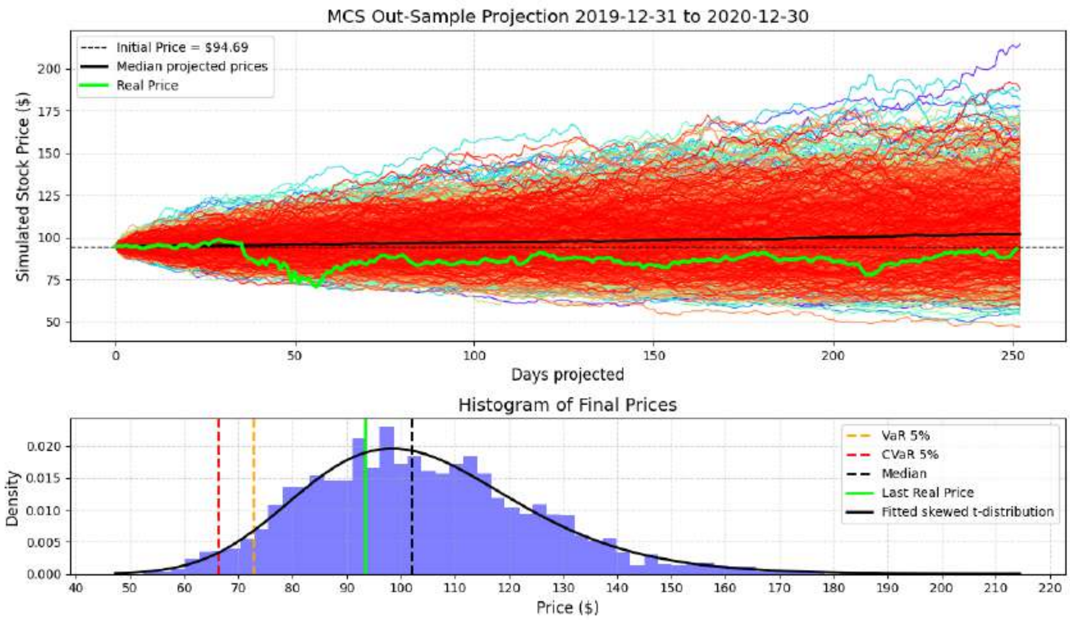

Results of backtesing (Target Price 01.2022)

Actual Price NVS (01.2022)	87$

Analysts	Target Price	Difference
Barclays	92$	6%
Bank of America	117$	34%
Deutsche Bank	105$	21%
Jefferies	93$	7%
JP Morgan Chase	102$	17%
Morgan Stanley	84$	-3%
Société Générale	86$	-1%
UBS	97$	11%
Vontobel	116$	33%
Mean	99$	14%

MCS on stock return	102$	17%
MCS on FCFF	98$	13%

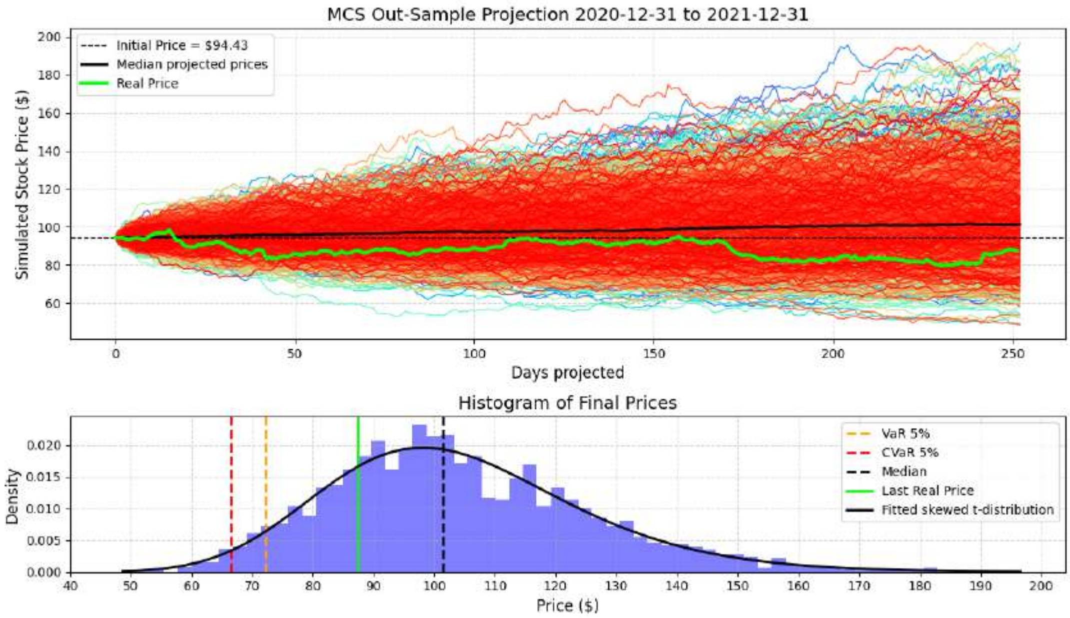

Results of backtesing (Target Price 01.2023)

Actual Price NVS (01.2023)	91$

Analysts	Target Price	Difference
Barclays	92$	1%
Bank of America	72$	-21%
Deutsche Bank	65$	-29%
Jefferies	63$	-31%
JP Morgan Chase	85$	-7%
Morgan Stanley	70$	-23%
Société Générale	90$	-1%
UBS	94$	3%
Vontobel	59$	-35%
Mean	77$	-16%

MCS on stock return	94$	3%
MCS on FCFF	92$	1%

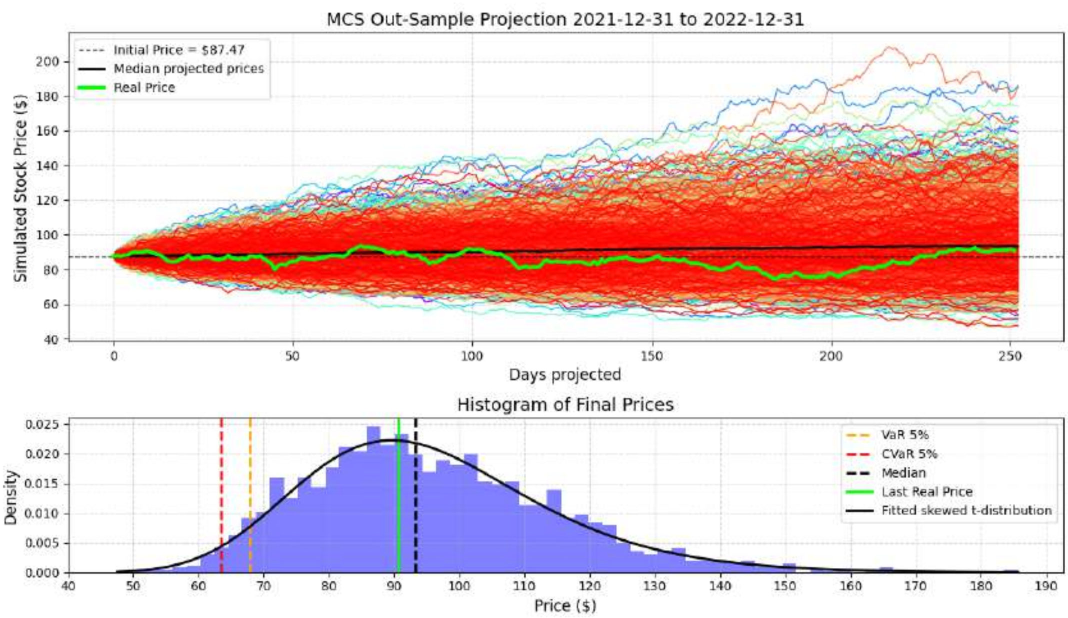

Results of backtesing (Target Price 01.2024)

Actual Price NVS (01.2024)	101$

Analysts	Target Price	Difference
Barclays	108$	7%
Bank of America	104$	3%
Deutsche Bank	98$	-3%
Jefferies	97$	-4%
JP Morgan Chase	110$	9%
Morgan Stanley	106$	5%
Société Générale	106$	5%
UBS	104$	3%
Vontobel	103$	2%
Mean	104$	3%

MCS on stock return	98$	-3%
MCS on FCFF	95$	-6%

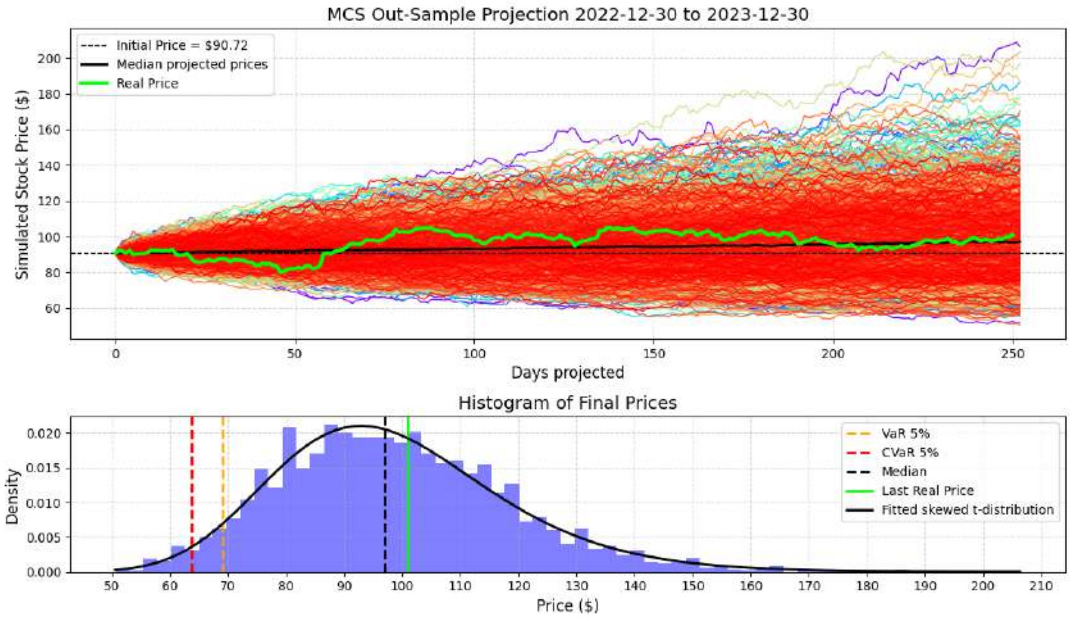

Complementary analysis of the VIX

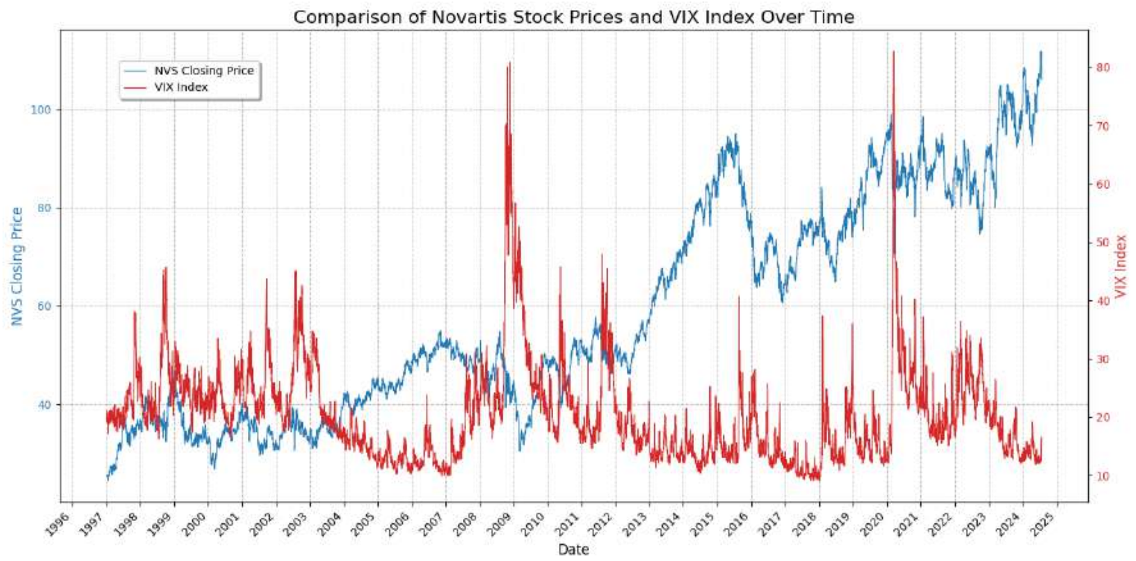

Figure 17: NVS & VIX

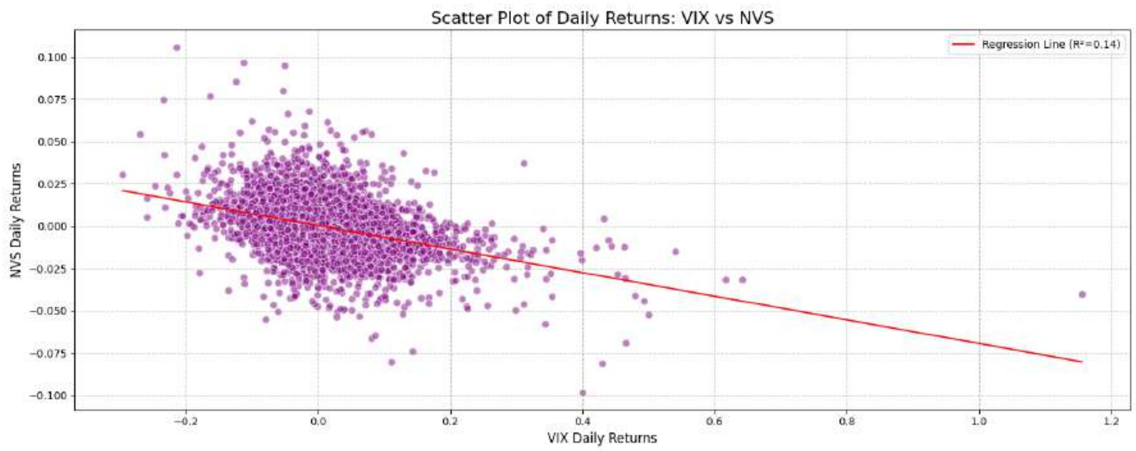

Figure 18: Correlation between NVS and VIX

Varing Moments

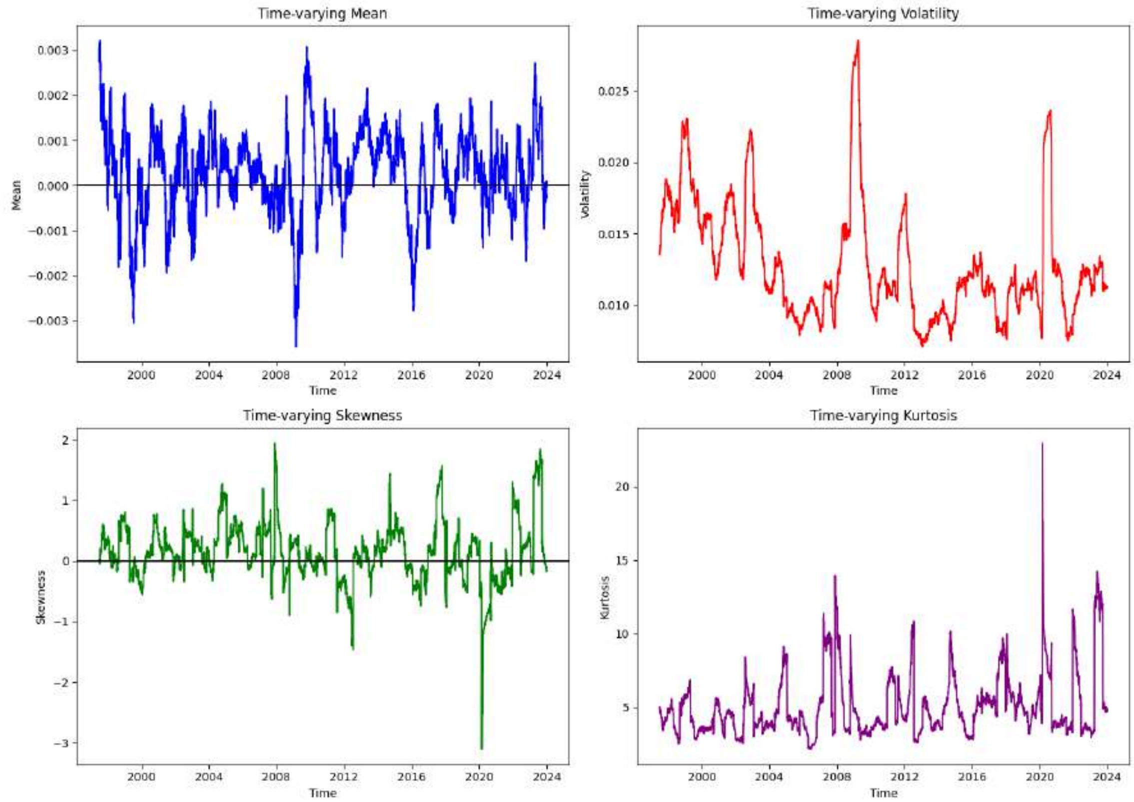

Figure 19: Varing Moments

List of Tables