Erasmus School of Economics

Thesis Quantitative Finance

July 17, 2023

Distributional Regression Forest applied on stock returns

Author

Pepijn Harsveldt

Student Number

498787

Supervisor:

dr. O. Kleen

Second Assessor:

dr. A . Camehl

Abstract

This paper proposes a novel approach using a Distributional Regression Forest (DRF) for probabilistic forecasting of stock returns. The DRF model goes beyond point estimation by leveraging the random forest framework to compute probabilistic properties, such as the conditional mean and standard deviation. By capturing non-linearities and interactions, the DRF model effectively incorporates changing market conditions and relevant information at each conditional node of the tree. Distributional forecasts of three American stocks (CVS, Microsoft, and Apple) in the American equity markets compares the DRF model with the GJR-GARCH model. The findings highlight the superiority of the DRF model in point forecasting and univariate distributional modeling over the GJR-GARCH model, while demonstrating comparable strength in multivariate distributional modeling.

Keywords: Distributional Regression Forest, GJR-GARCH, probabilistic forecasting, American equity markets

Contents

1 Introduction **1**

2 Literature Review **4**

2.1 Distributional Regression Forests 4

2.2 Copulae 5

3 Data **7**

3.1 Dependent Variables 7

3.2 Independent Variables 10

4 Methodology **12**

4.1 GJR-GARCH 13

4.2 Distributional Regression Forest 14

4.2.1 Tree Estimation 15

4.2.2 Tree Formation 16

4.2.3 Tree Aggregation 16

4.2.4 Prediction 17

4.3 Copula 17

4.3.1 Sklar's Theorem 18

4.3.2 Copulae 18

4.4 Model Evaluation 19

4.4.1 Mean Squared Error 19

4.4.2 Continuous Ranked Probability Score 20

4.4.3 Conditional Equal Predictive Ability 21

4.4.4 Kolmogrov-Smirnov test 21

5 Results **23**

5.1 Point forecast evaluation 23

5.2 Distribution forecast evaluation 25

5.3 Copula fit 27

5.4 Feature Importance 29

6 Conclusion **31**

A Appendix

A.1 ALE plots . . . . .

A.1.1 CVS mean . . . . .

A.1.2 CVS standard deviation . . . . .

A.1.3 Apple mean . . . . .

A.1.4 Apple standard deviation . . . . .

A.1.5 Microsoft mean . . . . .

A.1.6 Microsoft standard deviation . . . . .

Acronyms

ALE Accumulated Local Effect.

CDF Cumulative Distribution Function.

CEPA Conditional Equal Predictive Ability.

COVID Corona Virus Disease, SARS COV-2.

CRPS Continuous Ranked Probability Score.

DRF Distribution Regression Forest.

GAM Generalized Additive Model.

GAMLSS Generalized Additive Model Location Scale and Shape.

GARCH Generalized AutoRegressive Conditional Heteroskedasticity.

GJR Glosten-Jagannathan-Runkle.

JB Jarque-Bera.

KS Kolmogrov-Smirnov.

ML Machine Learning.

MSE Mean Squared Error.

OLS Ordinary Least Squares.

RF Random Forest.

1 Introduction

The study of financial time-series modelling and forecasting has captivated researchers' attention in econometrics, particularly within the realm of quantitative finance. This pursuit has yielded many theories and even garnered prestigious Nobel prizes, such as Sharpe (1964). The utilization of Machine Learning (ML) techniques in this domain is still relatively nascent. As a result, it is not being used to its full potential yet. This research aims to explore the use of the distributional regression forest (DRF), an innovative ML model, on stock returns in the American equities market. This research uses the holding period return of Apple, CVS and Microsoft, from 2 January 2003 to 30 December 2022. A rolling window of fifteen years is used to estimate the model.

The financial industry encompasses a wide range of stakeholders, including banks, consumers, regulators, and insurers, all of whom have a vested interest in stock returns. Stock returns play a crucial role in the financial decisions made by these parties. The practical application of ML in the field is gaining momentum, particularly among investment banks that rely on deep neural networks to generate point predictions for stock returns (Huang et al., 2020). It is important to note that the majority of ML models utilized in the financial sector, including deep neural networks, primarily offer point forecasts (Rapach and Zhou, 2020). However, there is an increasing trend towards probabilistic forecasting (Gneiting and Katzfuss, 2014), partly driven by recent global events such as the war in Ukraine (2022, 2023) and the COVID-19 pandemic (2020, 2021, 2022). Instead of solely estimating the conditional mean, probabilistic forecasting seeks to predict the characteristics of a specified distribution. In this study, the Gaussian distribution it the pre-defined distribution of interest.

Distributional forecasts introduce noise into their prediction as the forecast imposes both the conditional mean and standard deviation. The inclusion of the standard deviation in distributional forecasts provides several notable benefits. For example, the standard deviation quantifies uncertainty, enabling a more realistic understanding of forecasted outcomes. This acknowledgement of variability is crucial in financial decision-making, allowing stakeholders

to set appropriate expectations and assess the potential range of outcomes.

The research presented and discussed in this paper explores the potential application of the distributional regression forest model within the domain of stock returns forecasting. Traditionally employed in weather forecasting, the DRF model exhibits intriguing characteristics that suggest its applicability in financial forecasting. Notably, the DRF model bares similarities with the well-known random forest model, RF (Breiman, 2001), renowned for its robust predictive capabilities of stock return forecasting (Park et al., 2022) and is assumed to be a ML model (Lohrmann & Luukka, 2019). The DRF model, akin to the random forest model, adopts a tree-based structure. However, it distinguishes itself by estimating the conditional distribution at each node rather than focusing solely on the conditional mean. This feature sets the DRF model apart and holds promise for capturing the complete distributional properties of stock returns and is, therefore, a distributional forecast model.

This paper makes a contribution to the existing literature by applying the DRF model to stock returns, particularly within the context of financial forecasting. Notably, the combination of the DRF model with financial forecasting has yet to be explored in the literature. In order to assess the efficacy and relative performance of the DRF model, an essential step is to compare it with a competing model. In this research, an application of the Generalized Auto-Regressive Conditional Heteroskedasticity (GARCH) model is used as the benchmark. GARCH modela are widely utilised in financial econometrics to capture time-varying distributional properties, such as the conditional mean and standard deviation. This paper makes use of the GJR-GARCH model as its applied GARCH model (Glosten et al., 1993). By comparison, this research aims to provide insights into the strengths and weaknesses of the DRF model for stock return forecasting. This search for the performance of the DRF model leads to the research question:

"Can a Distributional Regression Forest outperform the GJR-GARCH model in modelling stock returns?"

To address the research question, defining the concept of out-performance is essential. This research is structured into several analyses, each focusing on different aspects. The first analysis centres on the point forecast properties, explicitly examining the conditional mean and variance separately. Subsequently, the research assesses the univariate distributional fit of the forecasts. This analysis examines how well the forecasted distributions capture the actual empirical distribution of the stock returns, this is split into several quantiles and the entire sample. The final analysis in this research revolves around the multivariate distributional fit of the forecasts. Multivariate analysis of stocks is the industry standard preferred by investors and financial institutions (Engle III & Sheppard, 2001) and is included in this research.

Copulae are employed to alleviate the computational burden associated with constructing multivariate distributions (Patton, 2013). Copulae provide a flexible framework for modelling the dependence structure among multiple variables, enabling the examination of the joint distribution of stock returns. By conducting these analyses, this research aims to provide a complete evaluation of the performance of the DRF model compared to the GJR-GARCH model.

The research has been structured into several parts. Firstly, an extensive literature review (Section 2) is conducted to provide context for the GJR-GARCH, DRF models and the copulae. Second, these models are compared to their respective rivals, evaluating their performance compared to other models. The third part of the research is the data section. This section begins with an overview of the dependent and independent variables, accompanied by their descriptive statistics. Following, the methodology section discusses the mathematical foundation of the models, estimation techniques, tests, and copula construction. The results section follows the methodology section, presenting the test statistics in an expanding format for comprehensive analysis. Finally, the research concludes with a discussion of the findings and recommendations for further research.

2 Literature Review

Economies are subject to fluctuations driven by various events and cyclical movements. As a result, this uncertainty plays a crucial role in the mispricing of financial assets. Traditional point forecast estimation methods often fail to capture this uncertainty explicitly. Upon this issue, distributional forecasts have gained prominence in financial forecasting (Gneiting & Katzfuss, 2014), as they provide a comprehensive view of the entire distribution of predictions rather than focusing solely on a single point estimate. By considering the entire distribution, these approaches offer valuable insights into the uncertainty surrounding predictions and facilitate more informed decision-making in the face of market fluctuations, ultimately contributing to an improved understanding of uncertainty dynamics and their implications in financial markets.

The pioneering work in distributional forecasting can be attributed to Koenker and Bassett (1978). In finance, several models have been employed for distributional forecasting, including the ARIMA (Box et al., 2015), Markov Chain Monte Carlo (Geyer, 1992), Gaussian Process (Williams and Rasmussen, 2006), Hidden Markov Model (Baum et al., 1972) and the GARCH model (Bollerslev, 1986). These models have played a significant role in capturing and understanding the distributional properties of financial variables, facilitating uncertainty into the forecasts in the realm of finance.

2.1 Distributional Regression Forests

This paper examines the suitability of employing a DRF model for modelling and predicting financial time series, particularly in the context of stock returns in American equity markets. Its structure is closely linked to the RF model, as discussed in Section 1. The tree-based structure of the DRF model serves as a statistical inference technique, primarily employed to enhance the smoothing of the target variables' responses, as elaborated in Section 4.2.3.

As discussed by Schlosser et al. (2019), the RF model can be classified within the class of Generalized Additive Models (GAM). GAMs are non-linear transformations of conventional

linear models, such as the ordinary least squares (OLS) model. On the contrary, the DRF model belongs to the class of Generalized Additive Models for Location Scale and Shape (GAMLSS), as highlighted by Schlosser et al. (2019). The DRF model adopts a distributional approach to the conventional GAM models by adding scale and location parameters. The classical framework of the DRF model assumes that the target variable cannot be negative, as Schlosser et al. (2019) explain. However, in this paper, alternative distributional assumptions are employed, and slight modifications to the likelihood function of the DRF model are made, as discussed in Section 4.2.1.

The performance of the DRF model is assessed by comparison with a highly implemented model, which in this research is the GJR-GARCH model. The GJR-GARCH model is widely recognised as a suitable and effective model for forecasting stock returns, as Alexander et al. (2021) and Xu and Zhu (2022) emphasise. The GJR-GARCH model builds upon the foundation of the GARCH (Bollerslev, 1986) and ARCH models (Engle, 1982). These models are based on the fundamental concept that modelling stock returns directly is challenging, while modelling variance is comparatively more feasible and informative. The GJR-GARCH model, known for its threshold parameters and consideration of multiple lags, is recognized for its flexibility and strong performance in capturing the dynamics of financial time series (Su et al., 2011).

2.2 Copulae

Portfolio managers, banks, regulators, and financial institutions continuously need to model and predict multiple stocks and financial instruments. Multivariate distributions are valuable for analyzing multiple stocks within a unified framework. As highlighted by Santos and Nogales (2012), multivariate models and distributions outperform their univariate counterparts, making them of significant interest to investors who prioritize comprehensive multivariate analysis over univariate analysis.

The construction of multivariate models can pose a computational burden, as noted by Patton (2013). However, Sklar's theorem (Sklar, 1959) offers a solution, as explained in 4.3.1.

This theorem establishes that multivariate distributions can be formed by utilizing copulae, which allows for constructing multivariate distributions from individual univariate distributions. Additionally, copula-based multivariate models offer flexibility in constructing the dependence structure of the distribution, not solely relying on the correlation matrix. Measures such as Spearman's rho (Spearman, 1987) or Kendall's tau (Kendall, 1938) can be employed to determine the underlying dependence structure. There are many classes within copulae, such as parametric and non-parametric[1].

The application of copulae in the financial domain is familiar and has been extensively utilised by various banks (Li, 2000). However, it is worth noting that the use of the Gaussian copula played a significant role in the financial crisis of 2008, earning it the notorious label of "The formula that killed Wall Street" (MacKenzie & Spears, 2014). It is essential to acknowledge that, at that time, tail-dependencies in copula modelling were not fully recognised or utilized to their full potential in mitigating systemic risks.

This paper incorporates four types of copulae, starting with the widely-used Gaussian copula, which offers high interpretability and finds extensive application in financial institutions. The second class of copulae employed in this study is the Vine copula, which utilizes a tree-based structure similar to a RF, providing greater flexibility in estimation, albeit at the cost of lower interpretability. Additionally, using Vine copulae helps alleviate the computational burden of copula computation.

The discussed Vine copulae include the Center Vine copula, which designates a center feature and conditions splits accordingly (Xiong and Cribben, 2022), the Regular Vine copula (Vine-R), which conditions splits on all previous features (Aas et al., 2009), and the Direct Vine copula, a simplified version that conditions on the previous split feature (Brechmann and Schepsmeier, 2013).

[1]For further detail in the different types of copulae please refer to Patton (2013)

3 Data

This section provides information on the databases and data sets utilized to address the research question outlined in Section 1. It specifies the sources employed and presents the essential characteristics of the dependent and independent variables. Consistent with the focus stated in Section 1, this paper examines American daily stock returns from 2 January 2003 to 30 December 2022. An important note is that the data is restricted to trading days, restricting to five trading days per week (public holidays, such as 4^{th} of July excluded). A month has approximately twenty-two trading days. The data analysis is structured into two main categories: dependent and independent variables, presented separately in this section.

3.1 Dependent Variables

The dependent variables in this study consist of stock returns from the American equity markets, as mentioned in Section 1. The rationale for focusing on American stock returns, rather than those of the European Union or other specific countries, is driven by the selection of independent variables and the overall stability of these stocks on a larger scale, as Solnik (1973) emphasised.

This research focuses on large US-based stocks: Microsoft, Apple, and CVS. The data for these stocks is obtained from the Center for Research in Securities Pricing (CRSP) and accessed through the Wharton Database platform. Accompanying, the triplet of Figure 1 provide a visual representation of the time series behaviour of each stock.

Figure 1: Time series of the individual stocks

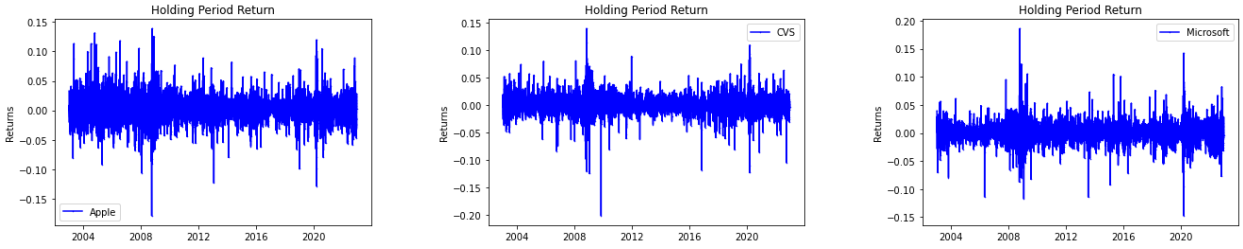

(a) Apple (b) CVS (c) Microsoft

Notes: Time-series of the different stocks, the y-axis differ among the different stocks. This is done in order to capture the entire time-series with extreme events.

This paper employs the holding period return as a measure of stock return. All stocks in this paper experienced stock splits in the time frame used. Using conventional methods to calculate stock returns would result in highly negative returns (due to a halving of prices)[2]. Therefore, the research adopts the holding period return, which accounts for stock splits and incorporates dividend spreads, providing a more accurate representation of an investor's daily return on their investment.

Apple exhibits greater volatility compared to Microsoft and CVS primarily due to its presence in the technology sector, which is known for its inherent volatility driven by rapid technological advancements, changing consumer preferences, and intense competition. While Microsoft is also classified as a technology company, its operations and diversification across multiple product lines provide stability and mitigate the impact of market movements. On the other hand, CVS operates in the healthcare sector, which is relatively less prone to the same level of volatility as the technology sector. Additionally, Apple's strong brand image, global market dominance, and the perception of high growth potential contribute to heightened market reactions and investor sentiment, further amplifying its stock return volatility. A thorough understanding of these factors is essential for investors to manage risk and make well-informed investment decisions effectively.

During the 2008 financial crisis and the COVID-19 pandemic (2020, 2021, 2022), Microsoft and CVS experienced the negative repercussions of the broader economic downturn and market uncertainties, leading to decreased consumer spending, disrupted supply chains, and challenging business environments in their respective sectors. However, as a healthcare-focused company, CVS faced additional challenges during the pandemic due to its direct involvement in addressing the healthcare crisis. While Apple was not immune to the impacts of these extreme events, its overall high volatility dampened the visibility of specific crisis effects on its stock performance. Notably, Apple exhibited the highest average stock returns among CVS, Microsoft, and itself, indicating the potential for greater returns. For a comprehensive

[2]The conventional ways of calculating stock returns is to either take the logarithmic difference or linear scaled difference.

understanding of the stocks' characteristics, including measures such as volatility and skewness, please refer to Table 3.1. The kurtosis is calculated by Fisher's formula (Fisher, 1930), indicating sample kurtosis.

Table 3.1: Descriptive Characteristics of the Dependent Variables

	Mean	**Std Dev**	**Skewness**	**Kurtosis**	**JB-statistic**[3]
Apple	0.150%	0.021	0.127	5.017	5294.335***
CVS	0.060%	0.017	-0.514	11.631	28601.521***
Microsoft	0.068%	0.017	0.229	9.889	20561.193***

Notes: This table contains the Descriptive Statistics of the stock returns for Microsoft, Apple and CVS. The JB-statistic is evaluated on three significance levels, * = 10%, ** = 5% and *** = 1%. The mean for CVS is the smallest combined with the smallest standard deviation. Apple has the closests fit to the normal distribution, but is still rejected being normally distributed. Especially CVS exhibits high Kurtosis and has negative skewness, indicating a right centered distribution.

An F-test was performed to compare the standard deviation between Microsoft and Apple, and between CVS and Apple. The results of both tests indicate no evidence to suggest a significant difference between the standard deviation of Apple's stock returns and the other stocks. Additionally, a t-test was conducted to compare the mean returns. The results indicate that Apple's mean return significantly differs from CVS's and Microsoft's. However, there is no significant difference in the mean returns between CVS and Microsoft.

Based on the mean and standard deviation, Table 3.1, the Microsoft and CVS stocks exhibit similar characteristics. The F and t-tests support this observation by indicating no significant difference in their mean returns. On the other hand, Apple stock demonstrates a distribution closest to normal compared to the other stocks. However, the assumption of normality is rejected based on the JB-statistic. All stocks, including Apple, Microsoft, and CVS, reject the null hypothesis of normality, indicating non-normal distribution. Moreover, all stocks exhibit leptokurtic distributions as their kurtosis values exceed 3, indicating heavy tails and a higher concentration of data around the mean. The accompanying histograms provide

[3]The Jarque-Bera statistic is used to determine the normality of a time-series. This is done by evaluating the Skewness and Kurtosis and converge them to a single output. For further information see Jarque and Bera (1980).

visual representations of the stock returns' underlying distributions.

Figure 2: Distribution of the dependent variables

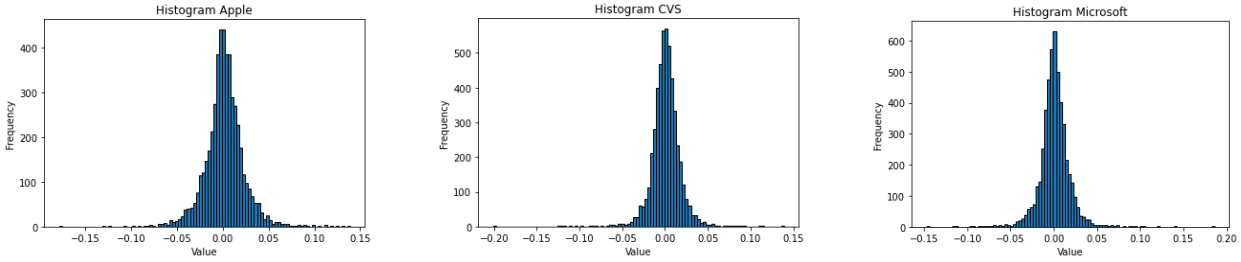

(a) Apple (b) CVS (c) Microsoft

Notes: This figure shows the distribution of the dependent variables for the different companies.

3.2 Independent Variables

In this section, the explanatory (independent) variables are examined. Consistent with the frequency of the dependent variables, the independent variables are measured daily. The data sources for the independent variables include the Center for Research in Securities Pricing (CRSP) and the (Archived) Fed Economic Research Division, (A)FRED. The independent variables included in this research are given in Table 3.2.

Table 3.2: Independent Variables used for the Distributional Regression Forest

Feature	**Description**
VIX	The VIX is the CBOE Volatility Index, providing insight into market volatility
S&P 500	Return series of the 500 biggest companies in the United States
Dow Jones Index	Value-weighted industrial average of the Dow Jones
NASDAQ	The Nasdaq Composite index containing approximately 3,000 stocks and is the first and foremost electronic trading platform
OIL	Crude Oil Prices: West Texas Intermediate (WTI) - Cushing, Oklahoma
GOLD	Producer Price Index by Industry: Gold Ore Mining: Gold Ores
SMB	Fama-French indicator, Small Minus Big
HML	Fama-French indicator, High Minus Low

Table 3.2 (continued)

Feature	Description
RMW	Fama-French indicator, Robust Minus Weak
CMA	Fama-French indicator, Conservative Minus Aggresive
DEXUSEU	U.S. Dollars to Euro Spot Exchange Rate
DGS1MO	Market Yield on U.S. Treasury Securities at 1-Month Constant Maturity, Quoted on an Investment Basis
BAMLH0A0HYM2	ICE BofA US High Yield Index Option-Adjusted Spread, Percent, Daily, Not Seasonally Adjusted
DPRIME	Exchange rate with EU
DHHNGSP	Henry Hub Natural Gas Spot Price, Dollars per Million BTU, Daily, Not Seasonally Adjusted
WILL5000INDFC	Wilshire 5000 Total Market Full Cap Index, Index, Daily, Not Seasonally Adjusted
USEPUINDXD	Economic Policy Uncertainty Index for United States, Index, Daily, Not Seasonally Adjusted
BAMLCC0A0CMTRIV	ICE BofA US Corporate Index Total Return Index Value, Index, Daily, Not Seasonally Adjusted
DPROPANEMBTX	Propane Prices: Mont Belvieu, Texas, Dollars per Gallon, Daily, Not Seasonally Adjusted
EMRATIO	Employment-Population Ratio, Percent, Monthly, Seasonally Adjusted
PCE	Personal Consumption Expenditures, Billions of Dollars, Monthly, Seasonally Adjusted Annual Rate

Notes: This table contains the independent variables used in order to fit the distributional regression forest. An importing note is that this table does not contain all explanatory variables, as the Skewness and Kurtosis is not taken into account.

Table 3.2 provides an overview of the external independent variables used in this study. However, this paper extends the external independent variables by incorporating lagged and transformed factors of the dependent variables as additional independent factors in the pre-

diction of the DRF model. Specifically, transformations such as skewness, kurtosis, and lagged volatility are considered, along with lagged returns on daily, weekly, and monthly intervals. The skewness factor is computed using Pearson's formula (Pearson, 1895), while the kurtosis factor is computed using Fisher's sample kurtosis formula (Fisher, 1930). By assuming a zero conditional mean for stock returns, Gropp (2004), the volatility of the return series can be estimated by its square. Furthermore, several dependent variables are linear transformed to return series, such as the S&P index and the D&J index.

4 Methodology

The Methodology section provides an in-depth understanding of the theoretical foundations of this research. For clarity and coherence, this section is structured into distinct components. Commencing with an exposition of the (GJR-)GARCH model. Secondly, a detailed description of the DRF model, followed by a description of the copula methodology. Subsequently, the test techniques employed in this study are displayed.

The GJR-GARCH and DRF model are estimated using the Gaussian distribution[4]. This distribution is chosen due to its relative resemblance to the market dynamics, enhancing the applicability and interpretability of our findings and industry standard. Moreover, to effectively capture the dynamic nature of the market, a rolling window approach is adopted, employing a fifteen years span. By opting for a rolling window instead of an expanding window, accountability for fundamental changes, such as the crisis in 2008 or the COVID-19 pandemic, that may occur over time are ensured, thereby fortifying the robustness of our analysis.

[4]The use of other distributions is possible, such as the skew-t distribution. However, the results between the Gaussian distribution and Skew-t distribution were negligible for the DRF model. Therefore, for the sake of simplicity, the Gaussian distribution is considered in this paper.

4.1 GJR-GARCH

GARCH models have emerged as a prominent tool for analyzing financial time-series, as discussed in Section 2. In the context of this research, the prediction horizon (denoted by t) covers the period from 2 January 2018 to 30 December 2022. This chosen time-frame offers a comprehensive scope to capture the intricate dynamics of the market. At the heart of the GARCH modelling framework lies the fundamental equation,

$$r_t = \mu_{t,GARCH} + \epsilon_{t,GARCH}. \tag{1}$$

The returns r_t, at Equation (1), are modelled by a time-dependent mean, denoted by $\mu_{t,GARCH}$ and an error term $\epsilon_{t,GARCH}$. GARCH models split the $\epsilon_{t,GARCH}$ in the following fashion,

$$\epsilon_{t,GARCH} = \sigma_{t,GARCH} z_t. \tag{2}$$

z_t represents a value derived from a predetermined distribution, in this study, z_t follows a normal distribution denoted as $z_t \sim \mathcal{N}(0, 1)$. While z_t is influenced by the assumed distribution, $\sigma_{t,GARCH}$ can be modelled independently. The modelling of $\sigma_{t,GARCH}$ lies at the core of GARCH models, which offer various specifications for its estimation. This paper employs the Glosten-Jagannathan-Runkle (GJR)-GARCH model proposed by Glosten et al. (1993). The choice of this model is motivated by its robustness in capturing and predicting financial time series, as discussed in Section 2. The GJR-GARCH model is specified as:

$$\sigma^2_{t,GARCH} = \omega + (\alpha + \gamma I_{t-1})\, \varepsilon^2_{t-1,GARCH} + \beta \sigma^2_{t-1,GARCH}. \tag{3}$$

ω is the conditional mean of the variance. In the GJR-GARCH model the $\epsilon^2_{t-1,GARCH}$ serves the $\sigma^2_{t,GARCH}$ in a dual way. At first by a direct output, where α serves as its factor parameter of influence. Furthermore, $\sigma^2_{t,GARCH}$ is influenced by the term $\gamma I_{t-1} \varepsilon^2_{t-1,GARCH}$, this term includes an indicator function specified as,

$$I_{t-1} := \begin{cases} 0 & \text{if } r_{t-1} \geq \mu_{GARCH} \\ 1 & \text{if } r_{t-1} < \mu_{GARCH} \end{cases}. \tag{4}$$

The indicator function determines whether the previous return was above (or below) the μ_{GARCH}, indicating a good (or bad) state. The GJR-GARCH model is an auto-regressive model, it uses its previous lag, $\sigma^2_{t-1,GARCH}$ and scales it by factor β. The parameter set $(\mu, \omega, \alpha, \gamma, \beta)$ is estimated simultaneously by log-likelihood estimation (for further explanation of log-likelihood estimation see Section 4.2.1).

4.2 Distributional Regression Forest

The structure of this section mirrors the methodology section of Schlosser et al. (2019). The first subsection focuses on the parameter estimation methodology for a specific tree, encompassing the estimation of the conditional mean and standard deviation associated with that tree. The second subsection delves into the creation of the trees, encompassing the split methods utilised and the depth of the trees. Lastly, the third subsection pertains to forming a smoothed tree, utilising the trees derived from the forest. These steps are visually represented in Figure 3. Figure 3.a illustrates the estimation of the mean and standard deviation, incorporating a tree split. Figure 3.b visually depicts a single tree split, highlighting the conditional mean and standard deviation estimation. Finally, Figure 3.c provides a visually smoothed forest. It is crucial to emphasise that these figures are purely illustrative and are not constructed using the data employed in this paper.

Figure 3: Tree estimation and Bagging method visual example

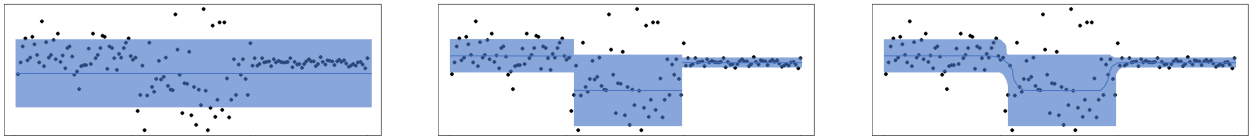

(a) Estimation of Mean and Standard Deviation for a Tree

(b) Partitioning and estimation of a tree

(c) Smoothed tree, upon full estimation of the DRF model

The DRF model is a deviation of a conventional random forest model (Meinshausen, 2006). As this paper heavily relies on the earlier work of Schlosser et al. (2019), the notation follows a similar pattern. Discussed in Section 3, there are three companies of interest, denoted by $k = 1, \ldots, 3$. The estimation of the smoothed tree is done by multiple trees. Each tree needs to be estimated and optimized, we denote a tree by $i = 1, \ldots, num - trees$. $num - trees$

is the number of trees used in order to create the forest for company k. This research follows the guideline set in Breiman (2001), 100 trees for large data-sets. Within each tree there are several nodes, each node is based on the optimal split and is always a multiple of 2^n, where n is the number of splits (maximized by the max depth of the tree). Each node has index $j = 1\ ,\ \dots\ ,\ num - of - nodes$. An important thing to note is that each tree is estimated over a different subset of the training set. This is random sampled with replacement, further discussed in Section 4.2.2. This results in different amount of nodes in each tree for a given stock.

4.2.1 Tree Estimation

This subsection discusses the estimation of $\mu_{t,DRF}$ and $\sigma_{t,DRF}$ in a node of a specific tree. The DRF model leverages on the framework set by the random forest model and therefore implicates node creation. Within each node the conditional mean and standard deviation is estimated. Due to the Gaussian distribution assumption[5], a log-likelihood parameter estimation method can be applied in order to estimate $\mu_{t,DRF}$ and $\sigma_{t,DRF}$ at each node. For notation convenience the combination of $\mu_{t,DRF}$ and $\sigma_{t,DRF}$ is defined as, $\theta_{t,DRF} = (\mu_{t,DRF}, \sigma_{t,DRF})$. The log-likelihood function for the DRF model[6] is defined as,

$$\ell(\theta_{t,DRF}; r) = \begin{cases} \log\left\{\frac{1}{\sigma_{t,DRF}} \cdot \phi\left(\frac{r-\mu_{t,DRF}}{\sigma_{t,DRF}}\right)\right\}, & \text{if } r \neq 0 \\ \log\left\{\Phi\left(\frac{-\mu_{t,DRF}}{\sigma_{t,DRF}}\right)\right\}, & \text{if } r = 0. \end{cases} \tag{5}$$

$\ell(\theta_{t,DRF}; r)$ has $\theta_{t,DRF}$ and the stock return series as its input, upon these inputs it gives its log-likelihood scores. ϕ is the Probability Distribution Function (PDF) of the normal distribution (neglecting the constant terms). This notation slightly differs from the notation in Schlosser et al. (2019). In Schlosser et al. (2019) the DRF model is applied on rain forecasts, which can not be negative. The optimal combination for θ is obtained by the Maximum Likelihood method,

[5]Other distributions can be used within the log-likelihood method for the estimating of the parameters of the DRF model. These potential deviations from the Gaussian distribution are stated in Schlosser et al. (2019).

[6]The log-likelihood function for the GJR-GARCH model has a similar structure, however, the mean and variance upon the estimation of the model is different.

$$\hat{\theta}_{t,DRF} = \underset{\theta_{t,DRF} \in \Theta_{t,DRF}}{\operatorname{argmax}} \sum_{l=1}^{a} \ell\left(\theta_{t,DRF}; r_i\right). \tag{6}$$

a is equal to the number of different combinations of $\theta_{t,DRF}$ are evaluated and $\Theta_{t,DRF}$ is the collection of all possible combinations within $\theta_{t,DRF}$ and is dependent on the structure of the tree. This method shows how to obtain the $\mu_{t,DRF}$ and $\sigma_{t,DRF}$ at a certain node for a tree.

4.2.2 Tree Formation

A tree in a DRF models consists of multiple nodes, consisting of segmented data at each node. An increasing segmentation of the data results equivalently in multiple node creations. The segmentation is done upon characteristics of the dependent variables and is conduction in a recursive way, this paper follows the frame-work set by Schlosser et al. (2019),

1. Given the current sample of observations (and splits), estimate (the set) $\theta_{t,DRF}$.
2. For each feature, compute the optimal split based on log-likelihood ceteris paribus of the other features.
3. Taking the optimal splits of each feature as known, evaluate the split with the highest improvement in log-likelihood.
4. Update the tree by incorporating the split and create the nodes accordingly.
5. Repeat step 1, 2, 3 and 4 unless there is no split available resulting in an improvement in log-likelihood score or the maximum-depth of the tree is reached.

The 'original' sample of observations is the input data given by the fifteen years rolling window. The maximum-depth is pre-determined by a K-fold cross-validation and set equal for all trees for the corresponding stocks. By splitting each tree, the conditional covariance is reduced, leading to an improvement in predictive performance Schlosser et al. (2019).

4.2.3 Tree Aggregation

In the realm of predictive modelling, a DRF comprises multiple trees that collectively form a forest. Each tree partitions the input data into several segments and computes $\theta_{t,DRF}$

accordingly. However, the discrepancies in $\theta_{t,DRF}$ across different segments can be abrupt and lacking smoothness. Bootstrapped aggregation (bagging) is employed as a solution for smoothening $\theta_{t,DRF}$. Bagging involves resampling from the "original" sample with replacement. For each resampled dataset, a tree is fitted and estimated. Bagging within the DRF framework facilitates the creation of a more refined and accurate model by mitigating abrupt transitions and enhancing the overall predictive capabilities.

The final result is made visual in Figure 3.c, where the boundaries of the mean and standard deviation are smoothed. Note that this graph is only for illustration purposes and is not based on real data. The optimal amount of trees is important, too few trees might evoke sharp transition periods and increased bias. In this situation an underfitting scenario emerges. However, when too much trees are used the tree is made to smooth and the standard deviation is too low. In this scenario we speak of overfitting.

4.2.4 Prediction

The prediction of $\theta_{prediction}$ is done in a similar approach as Schlosser et al. (2019). The 'k-nearest-neighbour-weights' approach is used to determine $\theta_{prediction}$[7]. The set of dependent features of the prediction (denoted as the information set) are evaluated and compared to segments of each tree. The (euclidean) distance between the segments of the fitted trees and the information set is obtained. Each node has its own distance to the information set and θ. The weight each node should receive upon creating $\theta_{prediction}$, is determined by the distance (smaller distance is better). The weights of the nodes are used to construct $\theta_{prediction}$. Where the number of nodes to consider, k, is determined by a k-fold cross-validation and in this research equal to fifteen.

4.3 Copula

This subsection introduces copulae and the underlying theories needed to apply them. The copula section commences with Sklar's theorem, which is the basis and the core of copulae.

[7]For further reading see Schlosser et al. (2019) or Lin and Jeon (2006)

Furthermore in Equation (9), the mapping of the marginal distribution into a joint distribution is given. This shows the formation of a multivariate distribution by a copula. After the Sklar theorem, the different copulae applied in this paper are presented.

4.3.1 Sklar's Theorem

The Sklar theorem lies the foundation of the aggregation of joint distributions upon several univariate distributions (see Section 2). As due to the use of Sklar's theorem the computational burden of computing joint (multivariate) distributions can be reduced significantly. The Sklar's theorem is defined as

$$\text{Let} \quad \mathbf{R} \equiv [r_1, \ldots, r_n]' \sim \mathbf{F}, \tag{7}$$

$$\text{with } r_k \quad \sim F_k \quad \text{then} \quad \exists \mathbf{C} \quad : \quad [0,1]^n \to [0,1], \tag{8}$$

$$\text{s.t.} \quad \mathbf{F}(\mathbf{y}) = \mathbf{C}\left(F_1\left(r_1\right), \ldots, F_n\left(r_n\right)\right) \forall \mathbf{y} \in \mathbb{R}^n. \tag{9}$$

Within this notation the $\mathbf{R}$ is the collection of univariate stock return series, which are distributed by distribution $\mathbf{F}$. The underlying marginal stock return series is given by F_i, where i is the i^{th} series. These series are defined by an underlying distribution F_i, whereas the cumulative distribution function (CDF) of these series are between zero and one.

The combination of univariate time-series combined into a joint distribution, has similar properties to the univariate time-series. Where the multiple univariate distributions CDF's are combined into a joint CDF, which is still between zero and one. In Equation (9) the mapping of the univariate distributions into a joint distribution is done by copula $\mathbf{C}$, here the copula is used to map the univariate distributions into a multivariate distribution.

4.3.2 Copulae

The application of copulae is widely accepted in several fields of science and practice, this has resulted in different type of copulae and specifications. In the literature section the use of copulae in financial forecasting is introduced and an introduction on the copulae used in this paper. The first copula of interest in this paper is the Gaussian copula, which does not

have an analytical formulation.

The copulae are split in two types in this paper, the Gaussian copula and three different Vine copulae. Vine copulae bare resemblance with random forests for their respective underlying frame-work (see Section 2). The underlying tree structure for the construction of the conditional copula as proposed by the Vine copulae, result in no analytical formula for the copula.

4.4 Model Evaluation

This paper answers the question of the appropriateness of the DRF model in several fields of the prediction, i.e. point forecast prediction, distributional fit and multivariate distributional fit. At first the Mean Squared Error (MSE) is determined, this method analysis the forecasted mean and standard deviation separately of each stock. Secondly, the CRPS scores are calculated, this statistic provides information concerning the distributional fit of the predictions. In order to provide information concerning the significance of the CRPS scores, a CEPA regression is made on the difference in CRPS scores for the DRF and GJR-GARCH model. Finally, a test is conducted on the multivariate fit, the Kolmogrov-Smirnov test.

4.4.1 Mean Squared Error

The MSE is a widely used metric in model comparison. The MSE measures the squared difference between the forecast output and the observed value of the time-series. The time denoted by t is the out-of-sample estimation time. Therefore, t=1 equals the first day of the forecasting, equal to 2 January 2018. Equation (10) hands the general framework of the MSE,

$$\text{MSE} = \frac{1}{T}\sum_{t=1}^{T}\left(r_t - \hat{r}_t\right)^2. \tag{10}$$

In this formula r_i is the actual value (commonly referred to as the observed value) and $\hat{r}_i$ is the predicted value. However, the forecast application in this research is not aimed at points but in conditional means and standard deviations. Resulting in two MSE tests. The first

MSE test is concerning the conditional mean, defined as,

$$\text{MSE}_{\text{mean}} = \frac{1}{T}\sum_{t=1}^{T}\left(r_t - \hat{\mu}_t\right)^2. \tag{11}$$

The conditional mean is the closed single value that probabilistic forecast measure can output. The second test is made upon the standard deviation. The difference between the standard deviation of the observed time-series and the predicted time-series is measured. This results in the following equation,

$$\text{MSE}_{\text{StdDev}} = \frac{1}{T}\sum_{t=1}^{T}\left(\sigma_t^2 - \hat{\sigma}_t^2\right)^2. \tag{12}$$

Here σ_t^2 is the variance at point t for stock k. The variance of a time series is not directly observable . Therefore, a proxy is taken, namely, the squared daily return r_t^2. Under the assumption that the mean return is equal to zero, the squared returns give a measure of variance. The obtained variance obtained from the model, either DRF or GJR-GARCH, is $\hat{\sigma}_t^2$.

4.4.2 Continuous Ranked Probability Score

The distributional fit is evaluated by a similar method used in Schlosser et al. (2019), this method is the Continuous Ranked Probability Score (CRPS). This score evaluates the entire distributional forecast by a single statistic in contrast to the MSE tests. In Gneiting and Raftery (2007), there are many reasons for the CRPS, the main argument for the use of CRPS is that it directly defines scoring rules in terms of predictive CDF. This paper follows the notation as described in Schlosser et al. (2019), which is a slight deviation from Gneiting and Raftery (2007)[8]. The notation,

$$\text{CRPS}(r, F) = \int_{-\infty}^{\infty}(F(z) - \mathbf{1}(r \leq z))^2 dz. \tag{13}$$

This notation makes use of the distribution $F(z)$, this is the CDF of the predicted distribution. This CDF is constructed on the mean and variance obtained as results from the estimation

[8]In Gneiting and Raftery (2007) the equation is accompanied with a negative sign, therefore negative scores are obtained and the most negative score indicates low predictive performance.

methods. The distribution is assumed to be Gaussian, however, this may differ when other distributions are evaluated. $\mathbf{1}(r \leq z)$ is the empirical distribution of the test point and is assumed to be the 'true' empirical distribution. The CRPS score measures the difference between the forecast distribution and the empirical distribution. Within this framework the test results vary among zero and plus infinity, where zero indicates a perfect fit with the empirical distribution.

4.4.3 Conditional Equal Predictive Ability

The Conditional Equal Predictive Ability (from now CEPA) test is developed by Giacomini and White (2006). This test is an evolved Diebold-Mariano (Diebold and Mariano, 2002) test. The CEPA test introduces an additional independent variable on the Diebold-Mariano test. The advantage of this test procedure is that a statement can be made on the drivers of the difference. As by the Diebold-Mariano test a regression is made on the difference and a constant, in the CEPA test the other independent variable provides influence. By this approach the constant explains the difference on the base of the unexplainable model difference.

This research uses the squared lagged returns as the added independent variable. As the GJR-GARCH model heavily relies on lagged variance for its estimate of the future variance and prediction. Under the market hypothesis of 'no free-lunch' the returns of the individual stocks are assumed to have a mean of zero. By applying this hypothesis, the squared return gives a measure for the variance of the time-series. The time-series is auto-correlated, therefore the analysis of the CEPA are done by analysing the Newey-West residuals.

4.4.4 Kolmogrov-Smirnov test

This paper follows the theoretical framework of Patton (2013), the Kolmogrov-Smirnov (KS) statistic is used in order to determine the Goodness-of-Fit of a copula. The KS-statistic makes use of the empirical copula, as the true copula is yet unknown, the empirical copula is defined as (Patton, 2013):

$$\hat{\mathbf{C}}_T(\mathbf{u}) \equiv \frac{1}{T}\sum_{t=1}^{T}\prod_{i=1}^{n}\mathbf{1}\left\{\hat{U}_{it} \leq u_i\right\}. \tag{14}$$

The $\mathbf{u}$ is the collection of univariate CDF's (collection of Equation 15). The indexation is as follows for this equation, the i determines the specific univariate distribution, in this research correspondents to a certain stock residual. The t is the total time of the forecast, in this research taken from the 2 January 2018 till 30 December 2022. The function ranks each observation by checking of the collection of univariate CDF's is larger compared to its peers, by this approach a joint CDF is formed. The univariate CDF's are created in the following manner,

$$\hat{F}_i(\varepsilon) \equiv \frac{1}{T+1}\sum_{t=1}^{T}\mathbf{1}\left\{\hat{\varepsilon}_{it} \leq \varepsilon\right\}. \tag{15}$$

The $\hat{F}_i(\varepsilon)$ is equivalent to the univariate CDF, $u_{i,t}$. As the copulae are formed on the residuals, this is the function of the residuals,

$$\hat{\varepsilon}_{it} \equiv \frac{r_{it} - \mu_i\left(\mathbf{Z}_{t-1};\hat{\alpha}\right)}{\sigma_i\left(\mathbf{Z}_{t-1};\hat{\alpha}\right)}. \tag{16}$$

As the DRF and GJR-GARCH model are both probabilistic forecast models, the $\mu_i\left(\mathbf{Z}_{t-1};\hat{\alpha}\right)$ and $\sigma_i\left(\mathbf{Z}_{t-1};\hat{\alpha}\right)$ are both forecasted. The $\mathbf{Z}_{t-1}$ is the prior knowledge up to the forecast that goes into the model and $\hat{\alpha}$ are the estimated parameters in order to compute the μ and σ. The estimated empirical distribution is used for the construction of the KS-statistic, which is defined by the following equation,

$$KS_C = \max_t \left|\mathbf{C}\left(\mathbf{U}_t;\hat{\theta}_T\right) - \hat{\mathbf{C}}_T\left(\mathbf{U}_t\right)\right|. \tag{17}$$

The $\mathbf{C}\left(\mathbf{U}_t;\hat{\theta}_T\right)$ is the estimated copula and $\hat{\mathbf{C}}_T\left(\mathbf{U}_t\right)$ is the empirical copula, which is constructed by Equation (14). The KS-statistic measures the distance between the empirical copula and the estimated copula and determines the level of fit by their relative distance. A small KS-statistic is assumed to be preferred over a large number as it indicates a stronger fit to the data. An important note is that this analysis is conducted on the residuals and the accompanied distributions of the residuals.

5 Results

The result section comprises multiple subsections encompassing distinct tests. Initially, the MSE scores to evaluate point forecast properties are computed, specifically the conditional mean and standard deviation. Subsequently, a test is conducted to assess the distributional fit for each stock individually. Lastly, the multivariate fit is analysed by the Kolmogrov-Smirnov test. This comprehensive evaluation yields valuable insights into the accuracy and performance of the DRF and GJR-GARCH models when predicting stock returns in the American equity market. Collectively, these findings enhance the understanding of the model's efficacy.

5.1 Point forecast evaluation

In this study, the DRF and GJR-GARCH model are utilised to generate predictions for the conditional mean and standard deviation at each time point within the out-of-sample period, spanning from 2 January 2018 to 30 December 2022. This section separately analyses the conditional mean and standard deviation, evaluating their performance using MSE scores. The equations employed for this analysis can be referred to in Section 4.4.1. The resulting test outcomes for the MSE tests are presented in Table 5.1.

Table 5.1: Mean Squared Error Scores, concentrated on the mean of the DRF and GJR-GARCH model

	Mean		Std Dev	
	DRF	**GJR-GARCH**	**DRF**	**GJR-GARCH**
Apple	0.081 %	0.082 %	0.001 %	0.001 %
CVS	0.026 %	0.064 %	0.001 %	0.002 %
Microsoft	0.076%	0.079%	0.001%	0.001%

Notes: For all stocks the DRF model outperforms the GJR-GARCH model. Microsoft imposes a large difference, which might be due to outliers and conditional heteroskedaticity. For CVS the difference is large and the conclusion can be drawn that for point forecasting the DRF model is more suitable.

Based on the statistics provided in Table 5.1, the DRF model outperforms the GJR-GARCH model for the prediction of the mean. This can be concluded from the fact that all MSE values are lower for the DRF model compared to the GJR-GARCH model. For the prediction of the

variance there is no clear favourite. However, note that the proxy of the observed variance is the squared daily holding period return, resulting in a small number. The triplet of Figure 4 shows the estimation of the mean of the GJR-GARCH and DRF model combined with the observed out-of-sample value of the stock return.

Figure 4: Mean stock returns and the observed stock value

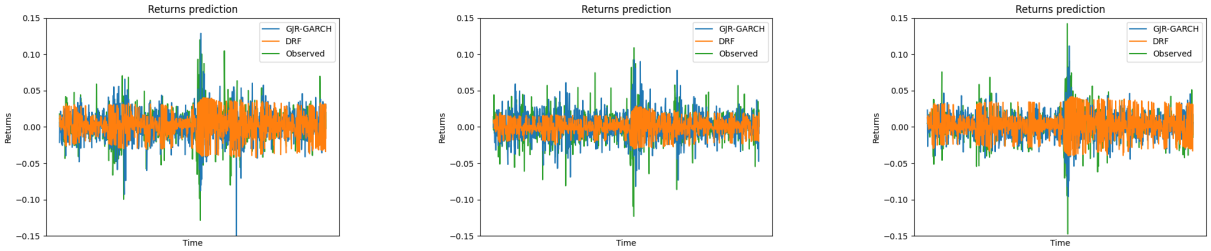

(a) Apple (b) CVS (c) Microsoft

In Figure 4, the orange line represents the predicted mean of the DRF model, equivalently the blue line for the GJR-GARCH model and green for the observed time-series. The GJR-GARCH seems to follow the observed time-series better. However, Figure 5 provides the absolute differences between the predicted mean and the observed values, this observation is rejected. Figure 5 showcases the obtained resulted by the MSE tests. In extreme events the GJR-GARCH model seem to provide a better fit, as the DRF model seems to be more conservative in its estimation of the mean.

Figure 5: Differences in observed stock value and estimated mean for the DRF and GJR-GARCH model

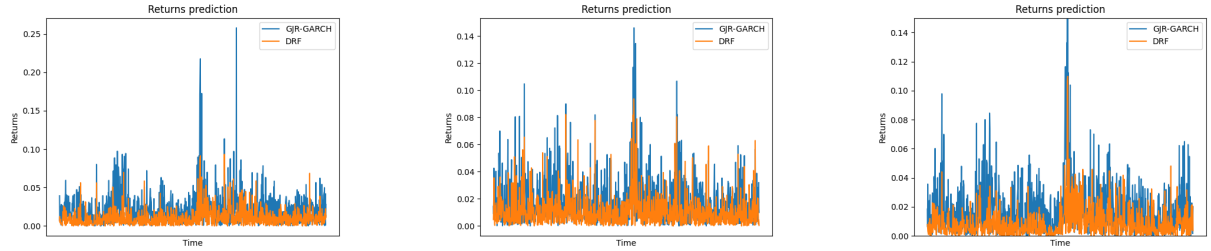

(a) Apple (b) CVS (c) Microsoft

5.2 Distribution forecast evaluation

The distributional forecast evaluation provides insight into the forecast's distributional fit upon the empirical distribution. This analysis (CRPS) is for each stock individually and has the conditional mean and standard deviation as inputs (Section 4.4.2). The test results of the CRPS analysis are presented in Table 5.2. As mentioned in Section 1, the testing procedure is divided into several sub-categories. In the analysis conducted by CEPA, the squared lagged return is used as a conditional variable for the significant difference analysis, as described in Section 4.4.3. Therefore, similar conditional sub-sections are employed in this analysis. However, the entire sample is also examined. The $33Q$, $MidQ$, and $67Q$ represent the sorted parts of the squared lagged returns. Specifically, 33Q refers to the smallest 33% of the sorted squared lagged returns for a specific stock.

Table 5.2: CRPS Scores of the Underlying Stocks, comparing the GJR-GARCH and the DRF model

	Full		33Q		MidQ		67Q	
	DRF	**GJR**	**DRF**	**GJR**	**DRF**	**GJR**	**DRF**	**GJR**
Apple	0.010^{*A}	0.015	0.010^{*A}	0.017	0.009^{*}	0.011	0.010^{*A}	0.016
CVS	0.010^{*A}	0.013	0.011^{*}	0.013	0.009^{*}	0.012	0.011^{*}	0.013
Microsoft	0.008^{*A}	0.013	0.008^{*A}	0.016	0.007^{A}	0.007	0.010^{*A}	0.015

*Notes: The CRPS scores are made for the full sample and several sub-samples. The sub-samples are based conditional the squared lag return as this measure is used in the analysis. The squared lagged return is a measure of volatility and is a key aspect of these models. The table showcases that the DRF model outperforms the GJR-GARCH model (denoted by GJR), as the CRPS score is lower. A CEPA test is conducted on the difference between the CRPS scores, a * denotes significant difference according to the CEPA test and ^{A} denotes that the squared lag return has significant explanatory power over the difference.*

The CEPA test has two outputs, namely, if the difference is significant and if the conditional feature (squared lagged returns) is significant. Whenever the conditional feature is significant, this tells us that the difference is either fully or partly driven by this feature. The significance of the conditional feature is denoted by ^{A}. Furthermore, when the difference is significant this is denoted by *. However, when the sign is both ^{*A}, this indicates that the

difference is significant even if the squared lagged returns are excluded. But, the squared lagged returns do have a significant effect, ceteris paribus. The CEPA indication is visible on the CRPS scores of the DRF model. For Microsoft the MidQ interval, is not significantly different in CRPS test scores.

The results from Table 5.2 provide a clear conclusion, namely, that in general the DRF model outperforms the GJR-GARCH model in the full set and nearly all sub-sets. For CVS only for the full set, the lagged squared returns are informative concerning the difference. For the sub-sets, conditioning on the squared lagged returns does not have a significant effect on the CRPS scores. In the full set the DRF models significantly outperforms the GJR-GARCH alternative. However, when analyzing the sub-sets this line does not fully concur.

Apple showcases pure superiority of the DRF model on the basis of CRPS test scores. Where all difference are significant even if the squared lagged returns do have a significant effect. When this effect is taken out, the differences remain significant. The 33Q subsection of the lagged squared return shows the largest difference in CRPS scores between the DRF model and GJR-GARCH model. This finding is in line with the previous finding that the DRF model is more conservative and estimates better in less volatile times. In the figure below a plot is provided, showing the CRPS scores of CVS over time from both the DRF and GJR-GARCH model.

Figure 6: Comparison of the DRF and GJR-GARCH model for the CRPS scores of CVS

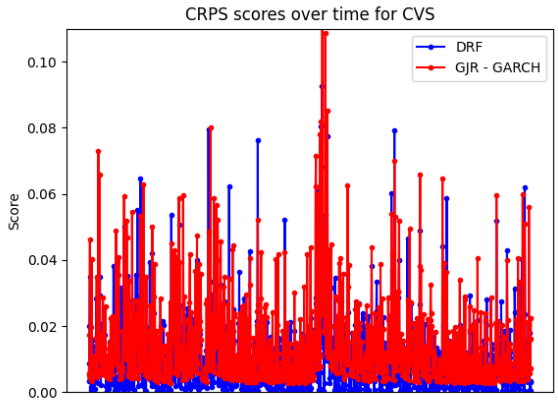

Notes: This figure is solely for the CVS stock. For all other stocks a similar pattern is followed. The blue line is for the DRF model and equivalently the red line is for the GJR-GARCH model. The base for the red line is higher compared to the blue line, indicating a higher standard level of mis-specification.

Figure 6 gives a visible interpretation of the results stated in Table 5.2, namely, the DRF model has a lower CRPS score on average and in extreme events, the DRF tends to have a less extreme result. This sections concludes on the superiority of the DRF model in forecasting the univariate distribution (assumed Gaussian) for the Apple, CVS and Microsoft stock.

5.3 Copula fit

The Kolmogrov-Smirnov statistic reflects the maximum distance between the 'true' empirical distribution and the fitted copula distribution (CDF). A small statistic represents a small difference among the true empirical distribution and therefore has a tighter fit on the data. Table 5.3 presents the Kolmogrov-Smirnov statistics for both the DRF as the GJR-GARCH model. Furthermore, the KS statistic is given for the full CDF and 95% of the CDF (the upper 5% of the CDF not included). The KS statistic is on average higher in the upper tail of the CDF (in this research) and has an effect on both estimation techniques. Therefore, an adjusted KS-test is conducted, where the upper tail of the CDF is not considered in computing the KS-statistic.

Table 5.3: Kolmogrov-Smirnov Scores of the copulae for the DRF and GJR-GARCH model

	Full		95%	
	DRF	**GJR**	**DRF**	**GJR**
Gaussian	0.092	0.212	0.028	0.091
Vine - Center	0.140	0.031	0.026	0.016
Vine - Regular	0.158	0.036	0.045	0.016
Vine - Direct	0.106	0.025	0.026	0.023

Notes: The KS statistics are showcased in this table. There is a division in either DRF or GJR-GARCH. Furthermore, the largest differences are for the DRF model in the upper tail of the CDF. As a result the KS statistics are presented without these tails.

In the full sample KS-test, the GJR-GARCH model outperforms the DRF model (except for the Gaussian copula). However, note that the KS-statistic is computed upon the empirical distribution of the residuals. There might be mis-specification present, which is especially visible in the tail of the CDF. The adjusted 95% sample of the KS-statistic bears a near similar conclusion, compared to the full sample. Namely, in the sub-sample the DRF model has a near similar KS-statistic, for the Vine-Direct copula (which indicate a near similar fit). The Gaussian copula shows a different results compared to the vine copulae, this is might be due to the underlying dependence structure. For the Gaussian copula, the DRF model presents a tighter fit.

The Kolmogrov-Smirnov statistic indicates whether the estimation models computes a strong fit to a multivariate distribution of the residuals (in this research). The 'true' distribution is yet unknown and therefore a proxy is taken, the empirical distribution. This might lead to mis-specification in tails. An important note is that the residuals are formed upon the observed time-series, the mean and standard deviation. The difference between the mean and observed time-series is smaller for the DRF model, according to the MSE scores. Therefore, the empirical distribution is different and has other shapes when compared to the GJR-GARCH model. Furthermore, the copula fit depends on the capability to capture the dependence structure and has its reflection on the KS-statistic. The GJR-GARCH model constructs its estimates based on previous estimates, the time-series of the three stocks is

correlated and therefore the estimates of the GJR-GARCH exhibits similar properties. The DRF model is a conditional model, resulting in a more complex dependence structure alongside the three stocks. This results in a higher KS-statistic. However, upon the KS-statistic there is no significant favourite between the DRF and GJR-GARCH model. This conclusion is partly driven by the 95% subset of KS statistics and the construction of the residuals derived by the model specific mean and variance predictions.

5.4 Feature Importance

The DRF model uses several input features, whereas not each feature is equally important. In order to increase the interpretability of the DRF model, Accumulated Local Effect (ALE) plots are made. ALE plots are preferred over partial dependence plots in this paper, due to the high correlation among the explanatory variables (Kleen and Tetereva, 2022). An ALE plot measures the change of the prediction due to a change of predictive features.

An important note, in this paper the DRF model outputs the conditional mean and standard deviation as its prediction, resulting in multiple ALE plots. The ALE plots in this paper follows the application set by Kleen and Tetereva (2022). The main difference in application is that not the difference in prediction is measured, but the difference in model output (here the μ_{DRF} and σ_{DRF})[9]. The difference between the ALE plots provided in Kleen and Tetereva (2022) and this paper is the aggregation. This paper evaluates each stock and characteristics separately, whereas in Kleen and Tetereva (2022) the aggregated ALE effect is investigated.

[9]For further understanding of the methodology and specific difference among conventional ALE plots and the ALE plots applied in this paper, please visit Kleen and Tetereva (2022)

Figure 7: ALE plots of Return-SP, Return-DJ and Return-SP for the influence on the mean of the prediction for Apple, CVS and Microsoft

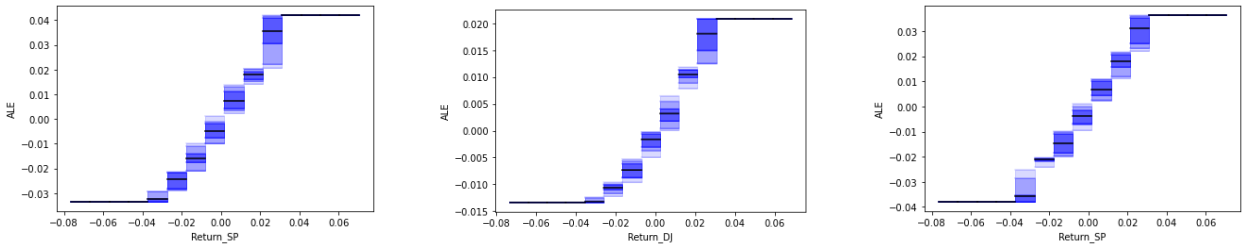

(a) Return-SP:Apple (b) Return-DJ: CVS (c) Return-SP:Microsoft

Similar to Kleen and Tetereva (2022) the parameter space is equally divided in fifteen bins. The return series of the S&P 500 and D&J are the most influential for all stocks. The black line determines the mean of the interval, whereas the dark blue line corresponds to the 50% interval. The medium blue surface corresponds to 80% and the light blue to the 95% interval of the effect on the prediction of the mean. A similar analysis is conducted upon the effect on the standard deviation. Yielding in the same set of predictors that are the most influential, the most influential features are displayed in this section. On the y-axis the ALE effect is presented, where the effect on the conditional mean or standard deviation is given conditional of the feature value (x-axis).

Figure 8: ALE plots of Return-SP, Return-DJ and Return-SP for the influence on the standard deviation of the prediction for Apple, CVS and Microsoft

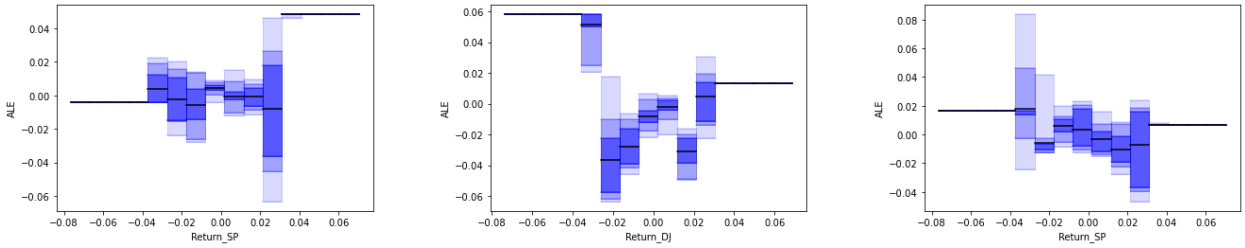

(a) Return-SP:Apple (b) Return-DJ: CVS (c) Return-SP:Microsoft

A clear result for the log-likelihood instead of conventional friendman-MSE optimisation is clear from the difference in ALE plots between the mean and standard deviation of the features in this section. The influence of a feature on the mean is not directly related to the

influence on the standard deviation, pressing on log-likelihood estimation. A large value on the y-axis of the ALE plot represents that if that feature value (x-axis) occurs, this will have a large effect on the estimation of the mean or standard deviation.

6 Conclusion

The conclusion section discusses the conclusion that can be drawn from the results and answers the research question of this paper: "Can a Distributional Regression Forest outperform the GJR-GARCH model in modelling stock returns?". As explained in the introduction, this question is directed in several sub-analysis.

The first analysis conducted is the MSE analysis, upon the test results in this subsection the answer is clear, the DRF model clearly out-performs the GJR-GARCH model. This result provides a valuable insight in the performance of the DRF model, that not only it performs well in distributional forecasting, i.e. predicting the conditional mean and standard deviation, but also in point forecasting it is strong. Furthermore, the GJR-GARCH model is specially designed to model the variance of the prediction. The DRF model has a similar value for the MSE compared to the GJR-GARCH model, indicating equal strength. The DRF is more conservative in its estimation for the mean and standard deviation compared to the GJR-GARCH model. Therefore, one might prefer the use of the GJR-GARCH model in extreme events and use the DRF model for less volatile moments. This insight is not tested in Schlosser et al. (2019) and provides information on the point forecast properties of the DRF model.

Univariate distributional fit is evaluated by the CRPS scores. The results discussed in Section 5.2 provide a clear result, the DRF outperforms the GJR-GARCH model. The CRPS score takes the conditional mean and standard deviation for the construction of the statistic. Upon this information, the conclusion can be drawn that the combination of distributional outputs of the DRF model is superior to the GJR-GARCH model. For Apple, the squared lagged returns deemed to be a good estimator for the difference in CRPS scores. Nevertheless, the

differences remained significant over all sub-sets and full set for Apple. The differences for CVS are significant in all sub-sets. But, the squared lagged returns do not have sufficient power to significantly explain the difference in CRPS scores (CEPA).

The last test conducted is the Kolmogrov-Smirnov test, this test evaluated the multivariate distributional fit of the residuals. The use of residuals is a linear transformation and the distribution is informative on the performance of the DRF model. There is no clear conclusion upon the KS-statistic, namely, for the full CDF the GJR-GARCH is better for the vine copulae. But, in the 95% adjusted sample there is no clear conclusion available upon the KS-statistics (difference in vine-direct is small). The Gaussian copula results in a contrary conclusion, where the DRF model out-performs the GJR-GARCH model. Potential mis-specification might be a issue combined with the complexity of the dependence structure. Furthermore, alongside different copulae the conclusion differs. More copulae and different type of copulae might change this result and conclusion upon.

ALE plots are implemented in order to increase the understanding of feature importance. The return series of the S&P 500 and D&J index are the most influential on the prediction of the conditional mean and standard deviation in this paper. Multiple return series would increase the accuracy of the DRF model, whereas the pure values of the indexes are less informative.

A DRF model has a higher run-time compared to the GJR-GARCH model. The fitting of trees and computing the predictions is a time consuming process, the optimisation of this is due for further research. A possible solution is the creation of a package, where the hyper-parameter tuning of the DRF model is included. Furthermore, the GJR-GARCH model solely depends on its assumed distribution and the lagged time-series returns. This research is limited to the Gaussian distribution, more distributions (Poisson, Weibull) might yield in different results. But, the choice of dependent variables for the DRF model is complexer. A broader data-set of explanatory variables might be beneficiary for further research. This research is limited to three US based stocks, more stocks in combination with different curren-

cies and geographical locations might enforce the conclusions upon this research and deepen the understanding of the DRF model. The use of copulae is limited to four different types, different copulae might give new insights in the multivariate distributional properties.

In order to answer the research question, the answer is yes. The DRF model provides a better prediction on average compared to the GJR-GARCH model for point and univariate distributional forecasting. However, the DRF model is more conservative and the results for the multivariate analysis does not provide a clear conclusion. The DRF model is a valuable tool in the field of financial forecasting.

References

Aas, K., Czado, C., Frigessi, A., & Bakken, H. (2009). Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, *44*(2), 182–198.

Alexander, C., Lazar, E., & Stanescu, S. (2021). Analytic moments for gjr-garch (1, 1) processes. *International Journal of Forecasting*, *37*(1), 105–124.

Baum, L. E., et al. (1972). An inequality and associated maximization technique in statistical estimation for probabilistic functions of markov processes. *Inequalities*, *3*(1), 1–8.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, *31*(3), 307–327.

Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time Series Analysis: Forecasting and Control*. John Wiley & Sons.

Brechmann, E. C., & Schepsmeier, U. (2013). Modeling dependence with c-and d-vine copulas: The r package cdvine. *Journal of Statistical Software*, *52*, 1–27.

Breiman, L. (2001). Random forests. *Machine Learning*, *45*, 5–32.

Diebold, F. X., & Mariano, R. S. (2002). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, *20*(1), 134–144.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the Econometric Society*, 987–1007.

Engle III, R. F., & Sheppard, K. (2001). Theoretical and empirical properties of dynamic conditional correlation multivariate garch.

Fisher, R. A. (1930). Moments and product moments of sampling distributions. *Proceedings of the London Mathematical Society*, *2*(1), 199–238.

Geyer, C. J. (1992). Practical markov chain monte carlo. *Statistical science*, 473–483.

Giacomini, R., & White, H. (2006). Tests of conditional predictive ability. *Econometrica*, *74*(6), 1545–1578.

Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, *48*(5), 1779–1801.

Gneiting, T., & Katzfuss, M. (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application*, *1*, 125–151.

Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, *102*(477), 359–378.

Gropp, J. (2004). Mean reversion of industry stock returns in the us, 1926–1998. *Journal of Empirical Finance*, *11*(4), 537–551.

Huang, J., Chai, J., & Cho, S. (2020). Deep learning in finance and banking: A literature review and classification. *Frontiers of Business Research in China*, *14*(1), 1–24.

Jarque, C. M., & Bera, A. K. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics letters*, *6*(3), 255–259.

Kleen, O., & Tetereva, A. (2022). A forest full of risk forecasts for managing volatility. *Available at SSRN*.

Koenker, R., & Bassett, G. (1978). Regression quantiles. *Econometrica*, *46*(1), 33–50.

Li, D. X. (2000). On default correlation: A copula function approach. *The Journal of Fixed Income*, *9*(4), 43–54.

Lin, Y., & Jeon, Y. (2006). Random forests and adaptive nearest neighbors. *Journal of the American Statistical Association*, *101*(474), 578–590.

Lohrmann, C., & Luukka, P. (2019). Classification of intraday S&P 500 returns with a random forest. *International Journal of Forecasting*, *35*(1), 390–407.

MacKenzie, D., & Spears, T. (2014). 'the formula that killed wall street': The gaussian copula and modelling practices in investment banking. *Social Studies of Science*, *44*(3), 393–417.

Meinshausen, N. (2006). Quantile regression forests. *Journal of Machine Learning Research*, *7*, 983–999.

Park, H. J., Kim, Y., & Kim, H. Y. (2022). Stock market forecasting using a multi-task approach integrating long short-term memory and the random forest framework. *Applied Soft Computing*, *114*, 108106.

Patton, A. (2013). Copula methods for forecasting multivariate time series. *Handbook of Economic Forecasting*, *2*, 899–960.

Pearson, K. (1895). Skew variation in homogeneous material.(contribution to thé mathematical theory of evolution ii). *Phil. Trans. of the Royal Society*, *186*, 43–414.

Rapach, D. E., & Zhou, G. (2020). Time-series and cross-sectional stock return forecasting: New machine learning methods. *Machine learning for Asset Management: New developments and financial applications*, 1–33.

Santos, A., & Nogales, F. (2012). Comparing univariate and multivariate models to forecast portfolio value-at-risk. *Journal of Financial Econometrics*, *11*.

Schlosser, L., Hothorn, T., Stauffer, R., & Zeileis, A. (2019). Distributional regression forests for probabilistic precipitation forecasting in complex terrain. *The Annals of Applied Statistics*, *13*(3).

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, *19*(3), 425–442.

Sklar, M. (1959). Fonctions de repartition an dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris*, *8*, 229–231.

Solnik, B. H. (1973). Note on the validity of the random walk for european stock prices. *The Journal of Finance*, *28*(5), 1151–1159.

Su, Y., Huang, H., & Lin, Y. (2011). Gjr-garch model in value-at-risk of financial holdings. *Applied Financial Economics*, *21*(24), 1819–1829.

Williams, C. K., & Rasmussen, C. E. (2006). *Gaussian processes for machine learning* (Vol. 2). MIT press Cambridge, MA.

Xiong, X., & Cribben, I. (2022). Beyond linear dynamic functional connectivity: A vine copula change point model. *Journal of Computational and Graphical Statistics*, 1–20.

Xu, Y., & Zhu, F. (2022). A new gjr-garch model for -valued time series. *Journal of Time Series Analysis*, *43*(3), 490–500.

A Appendix

A.1 ALE plots

A.1.1 CVS mean

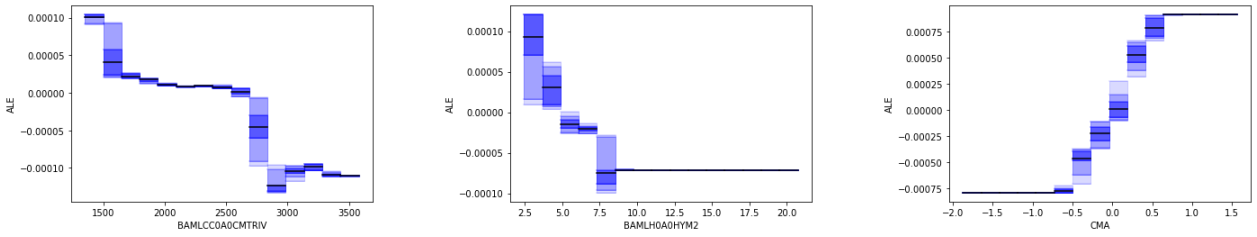

(a) BAMLCC0A0CMTRIV (b) BAMLH0A0HYM2 (c) CMA

Figure 9: ALE plots of BAMLCC0A0CMTRIV, BAMLH0A0HYM2 and CMA, dependent on the mean of CVS

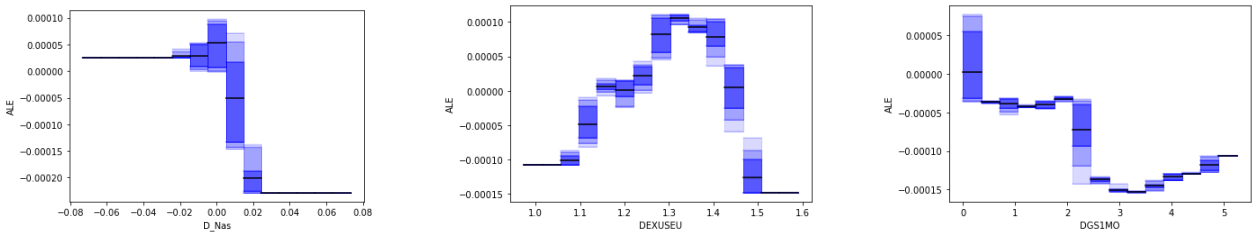

(a) D-Nas (b) DEXUSEU (c) DGS1MO

Figure 10: ALE plots of D-Nas, DEXUSEU and DGS1MO, dependent on the mean of CVS

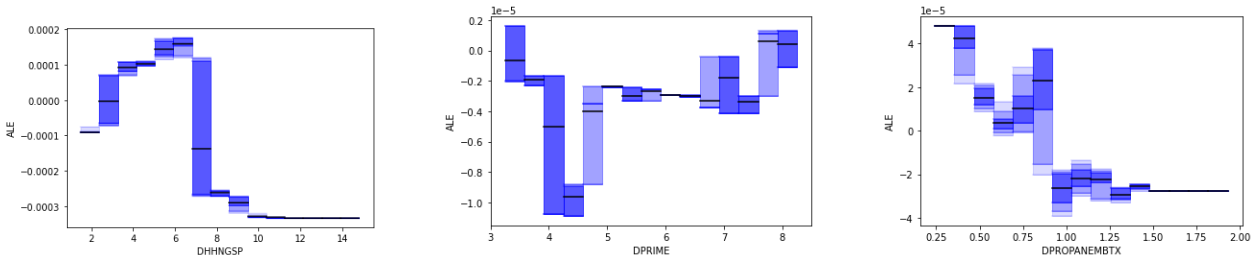

(a) DHHNGSP (b) DPRIME (c) DPROPANEMBTX

Figure 11: ALE plots of DHHNGSP, DPRIME and DPROPANEMBTX, dependent on the mean of CVS

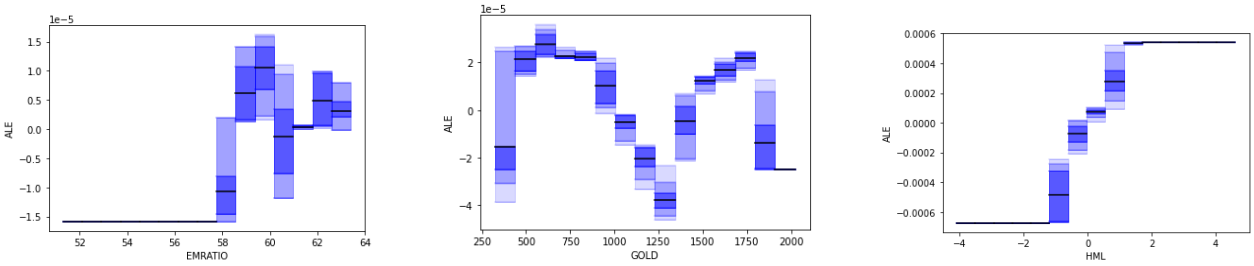

(a) EMRATIO (b) GOLD (c) HML

Figure 12: ALE plots of EMRATIO, GOLD and HML, dependent on the mean of CVS

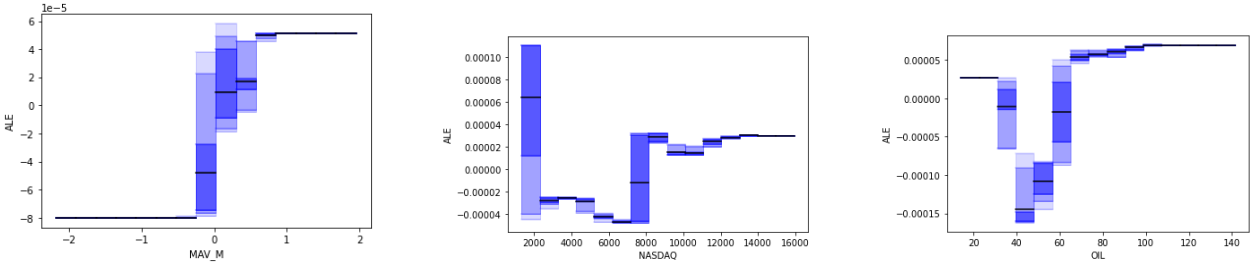

(a) MAV-M (b) NASDAQ (c) OIL

Figure 13: ALE plots of MAV-M, NASDAQ and OIL, dependent on the mean of CVS

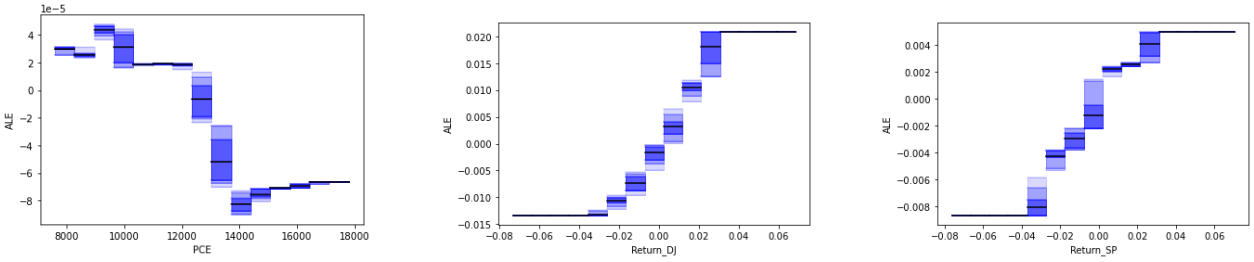

(a) PCE (b) Return-DJ (c) Return-SP

Figure 14: ALE plots of PCE, Return-DJ and Return-SP, dependent on the mean of CVS

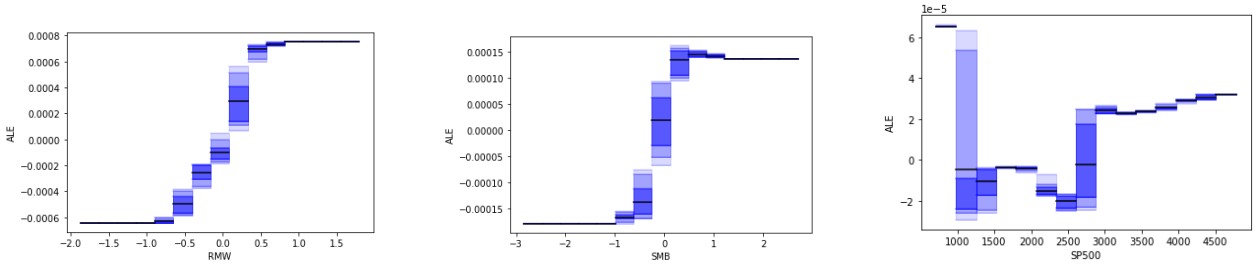

(a) RMW (b) SMB (c) SP500

Figure 15: ALE plots of RMW, SMB and SP500, dependent on the mean of CVS

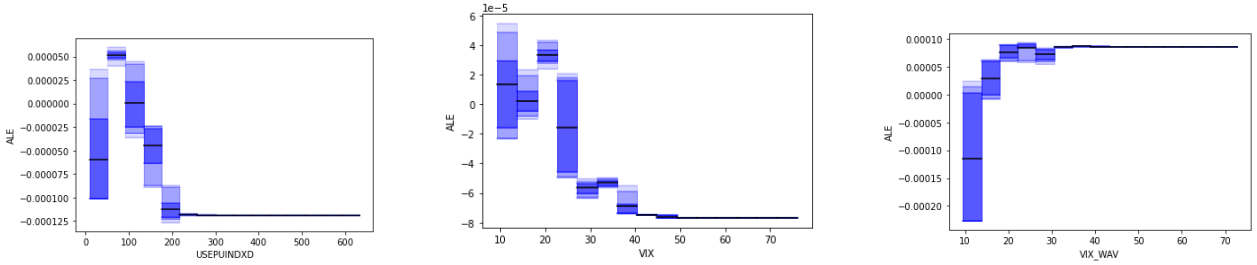

(a) USEPUINDXD (b) VIX (c) VIX-WAV

Figure 16: ALE plots of USEPUINDXD, VIX and VIX-WAV, dependent on the mean of CVS

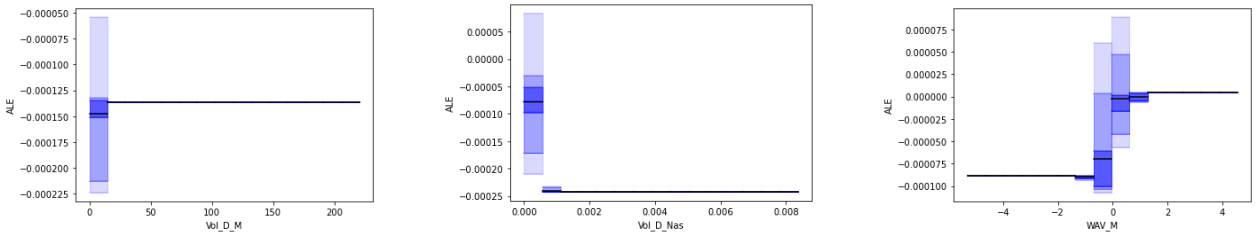

(a) Vol-D-M (b) Vol-D-Nas (c) WAV-M

Figure 17: ALE plots of Vol-D-M, Vol-D-Nas and WAV-M, dependent on the mean of CVS

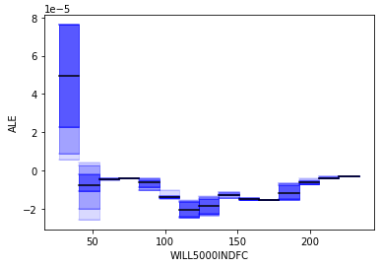

(a) WILL5000INDFC

Figure 18: ALE plots of WILL5000INDFC, dependent on the mean of CVS

A.1.2 CVS standard deviation

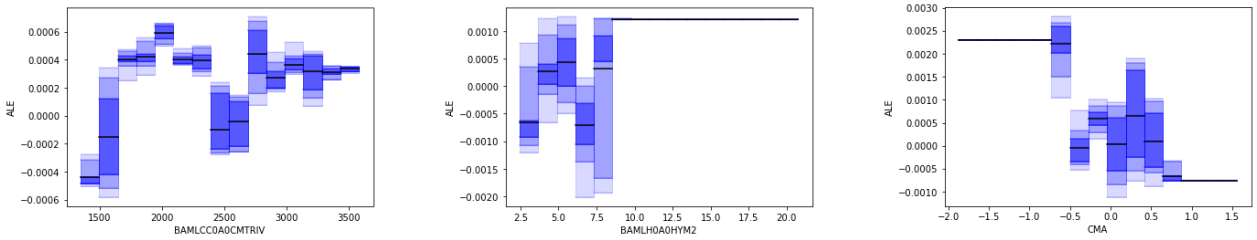

(a) BAMLCC0A0CMTRIV (b) BAMLH0A0HYM2 (c) CMA

Figure 19: ALE plots of BAMLCC0A0CMTRIV, BAMLH0A0HYM2 and CMA, dependent on the standard deviation of CVS

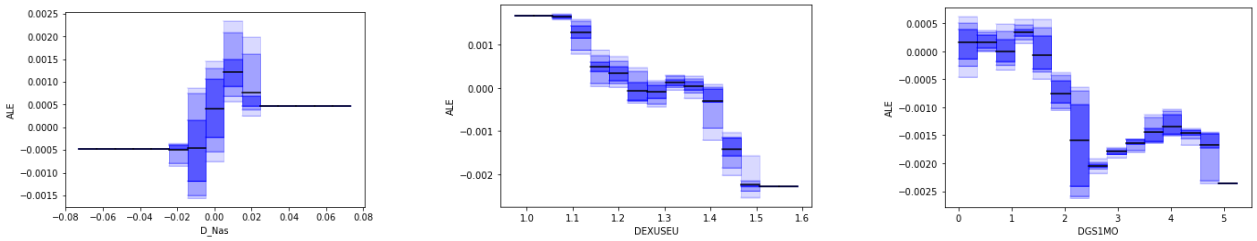

(a) D-Nas (b) DEXUSEU (c) DGS1MO

Figure 20: ALE plots of D-Nas, DEXUSEU and DGS1MO, dependent on the standard deviation of CVS

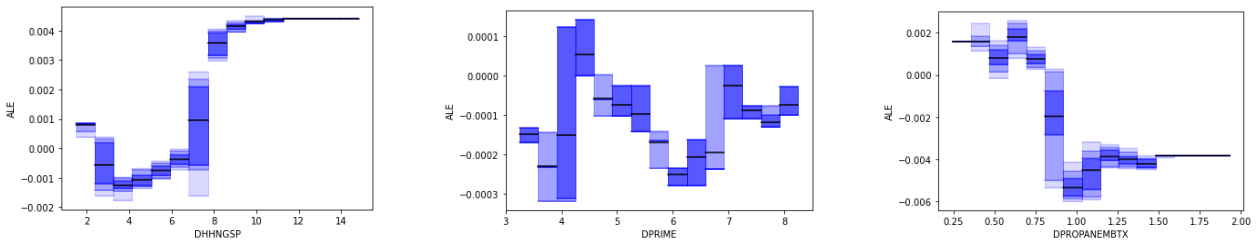

(a) DHHNGSP (b) DPRIME (c) DPROPANEMBTX

Figure 21: ALE plots of DHHNGSP, DPRIME and DPROPANEMBTX, dependent on the standard deviation of CVS

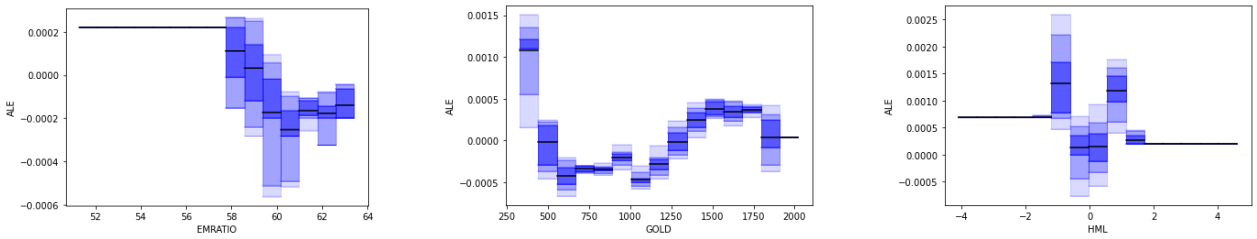

(a) EMRATIO (b) GOLD (c) HML

Figure 22: ALE plots of EMRATIO, GOLD and HML, dependent on the standard deviation of CVS

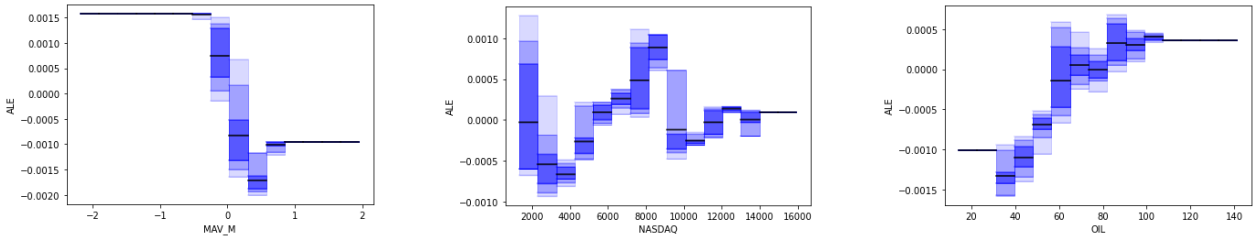

(a) MAV-M (b) NASDAQ (c) OIL

Figure 23: ALE plots of MAV-M, NASDAQ and OIL, dependent on the standard deviation of CVS

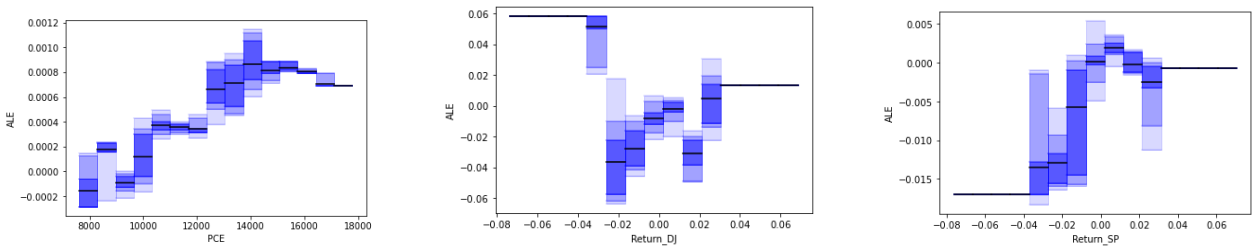

(a) PCE (b) Return-DJ (c) Return-SP

Figure 24: ALE plots of PCE, Return-DJ and Return-SP, dependent on the standard deviation of CVS

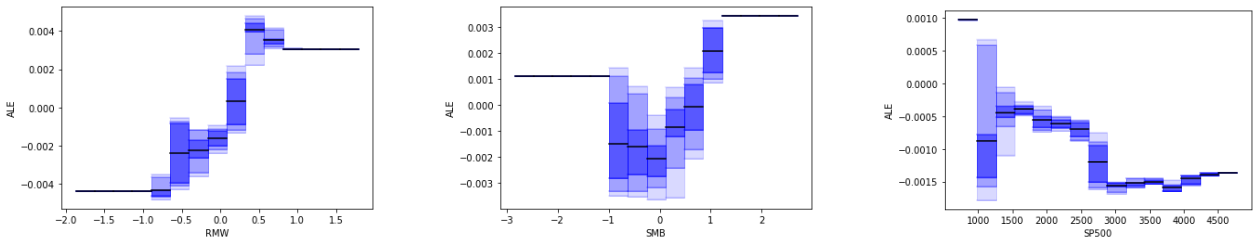

(a) RMW (b) SMB (c) SP500

Figure 25: ALE plots of RMW, SMB and SP500, dependent on the standard deviation of CVS

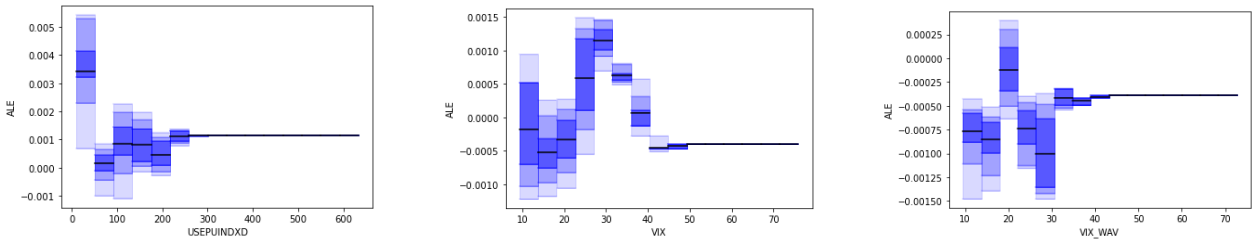

(a) USEPUINDXD (b) VIX (c) VIX-WAV

Figure 26: ALE plots of USEPUINDXD, VIX and VIX-WAV, dependent on the standard deviation of CVS

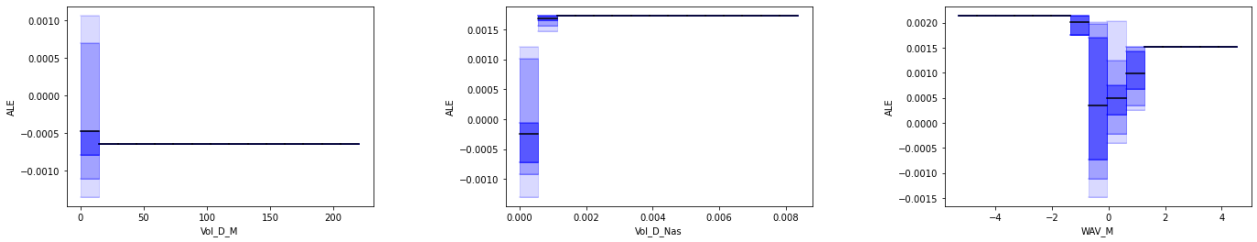

(a) Vol-D-M (b) Vol-D-Nas (c) WAV-M

Figure 27: ALE plots of Vol-D-M, Vol-D-Nas and WAV-M, dependent on the standard deviation of CVS

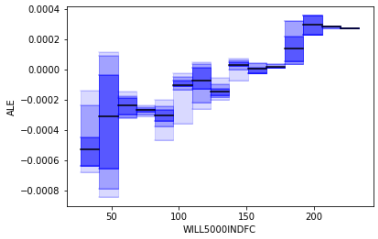

(a) WILL5000INDFC

Figure 28: ALE plots of WILL5000INDFC, dependent on the standard deviation of CVS

A.1.3 Apple mean

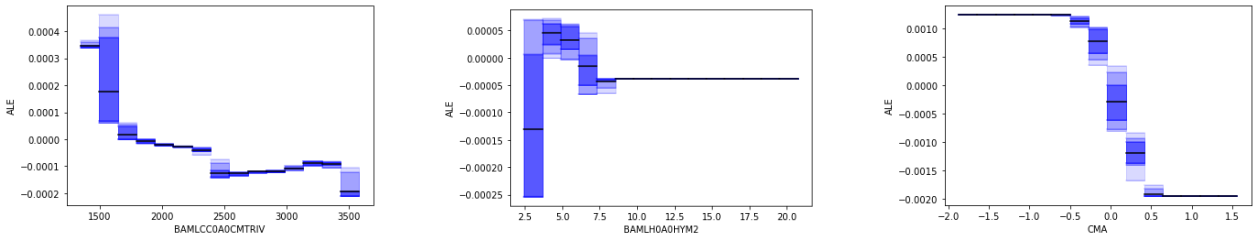

(a) BAMLCC0A0CMTRIV (b) BAMLH0A0HYM2 (c) CMA

Figure 29: ALE plots of BAMLCC0A0CMTRIV, BAMLH0A0HYM2 and CMA, dependent on the mean of Apple

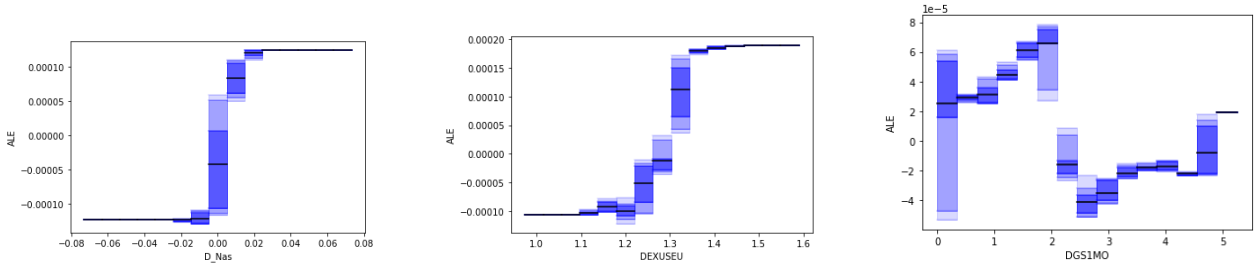

(a) D-Nas (b) DEXUSEU (c) DGS1MO

Figure 30: ALE plots of D-Nas, DEXUSEU and DGS1MO, dependent on the mean of Apple

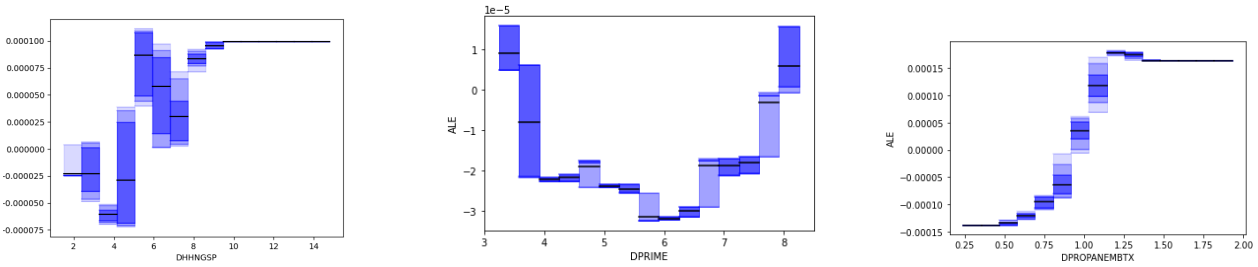

(a) DHHNGSP (b) DPRIME (c) DPROPANEMBTX

Figure 31: ALE plots of DHHNGSP, DPRIME and DPROPANEMBTX, dependent on the mean of Apple

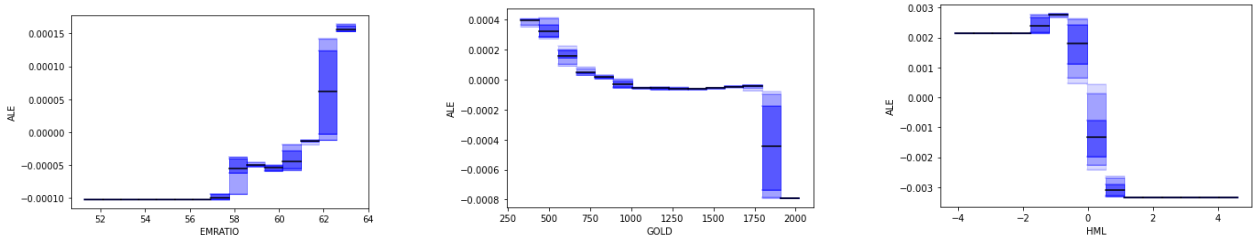

(a) EMRATIO (b) GOLD (c) HML

Figure 32: ALE plots of EMRATIO, GOLD and HML, dependent on the mean of Apple

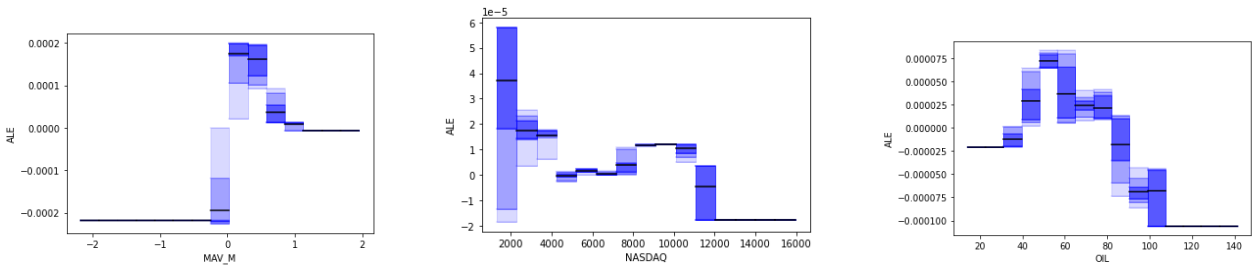

(a) MAV-M (b) NASDAQ (c) OIL

Figure 33: ALE plots of MAV-M, NASDAQ and OIL, dependent on the mean of Apple

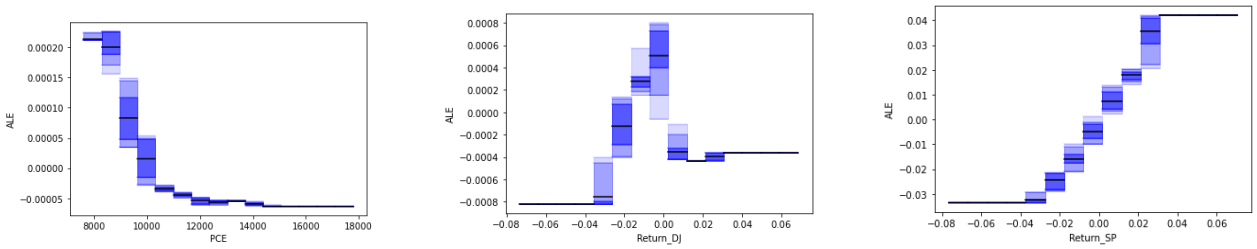

(a) PCE (b) Return-DJ (c) Return-SP

Figure 34: ALE plots of PCE, Return-DJ and Return-SP, dependent on the mean of Apple

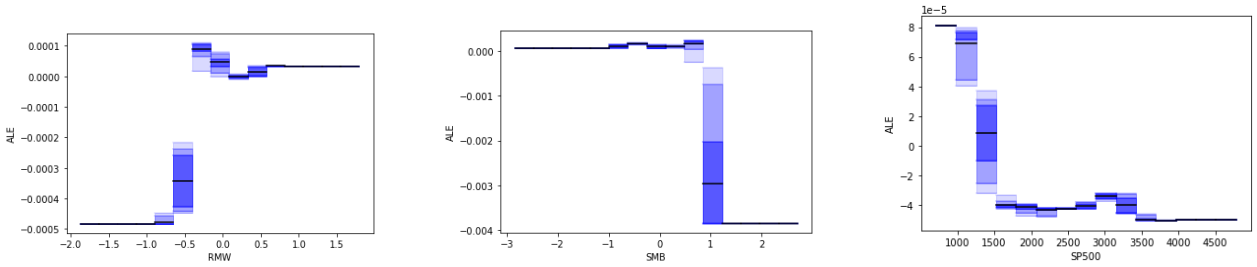

(a) RMW (b) SMB (c) SP500

Figure 35: ALE plots of RMW, SMB and SP500, dependent on the mean of Apple

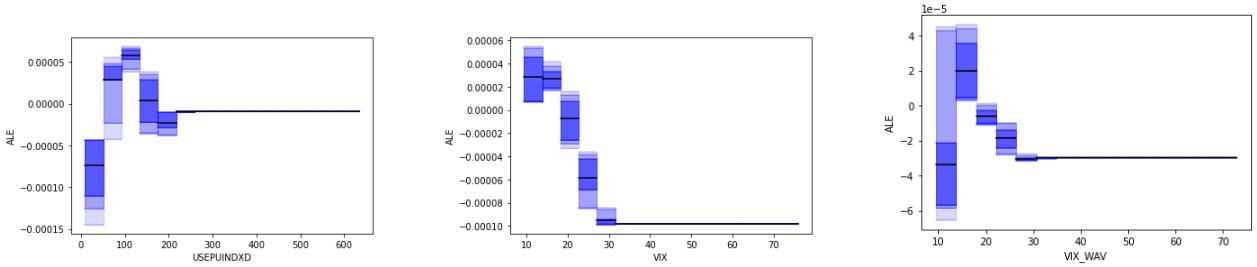

(a) USEPUINDXD (b) VIX (c) VIX-WAV

Figure 36: ALE plots of USEPUINDXD, VIX and VIX-WAV, dependent on the mean of Apple

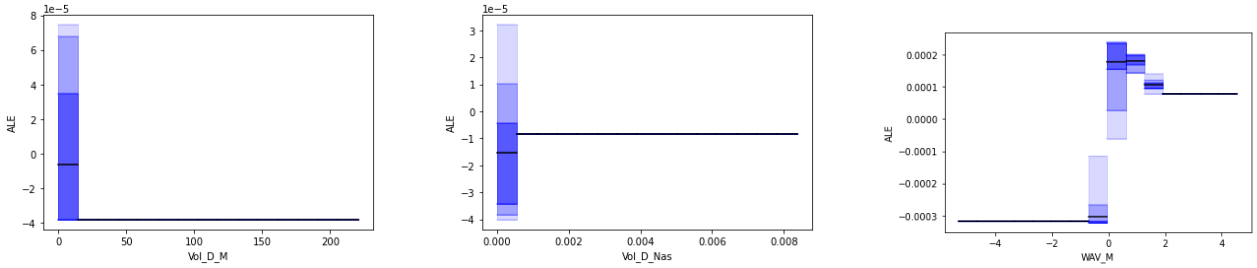

(a) Vol-D-M (b) Vol-D-Nas (c) WAV-M

Figure 37: ALE plots of Vol-D-M, Vol-D-Nas and WAV-M, dependent on the mean of Apple

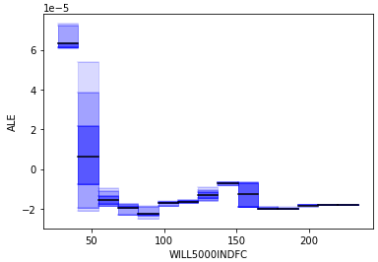

(a) WILL5000INDFC

Figure 38: ALE plots of WILL5000INDFC, dependent on the mean of Apple

A.1.4 Apple standard deviation

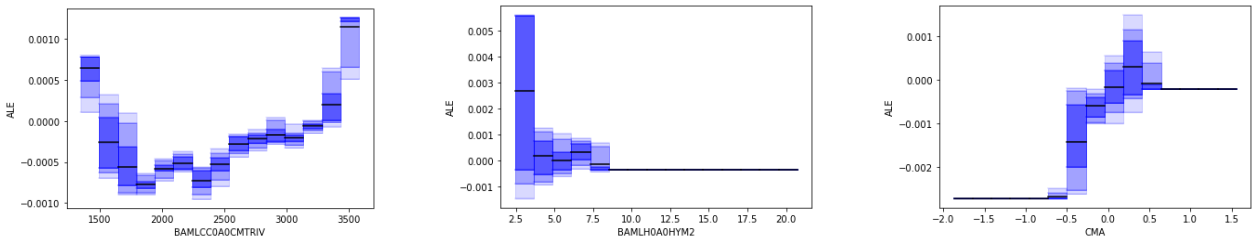

(a) BAMLCC0A0CMTRIV (b) BAMLH0A0HYM2 (c) CMA

Figure 39: ALE plots of BAMLCC0A0CMTRIV, BAMLH0A0HYM2 and CMA, dependent on the standard deviation of Apple

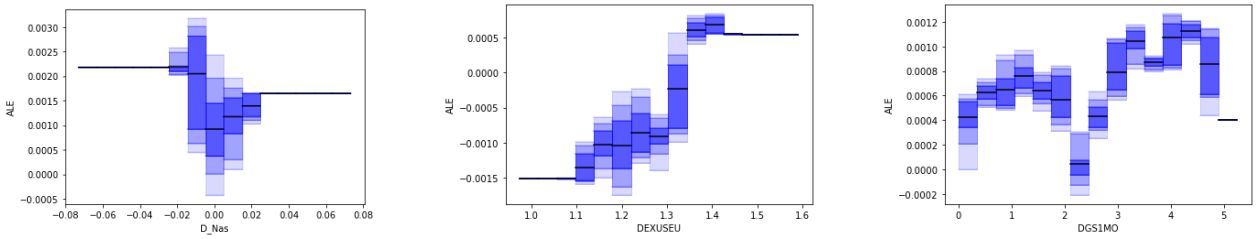

(a) D-Nas (b) DEXUSEU (c) DGS1MO

Figure 40: ALE plots of D-Nas, DEXUSEU and DGS1MO, dependent on the standard deviation of Apple

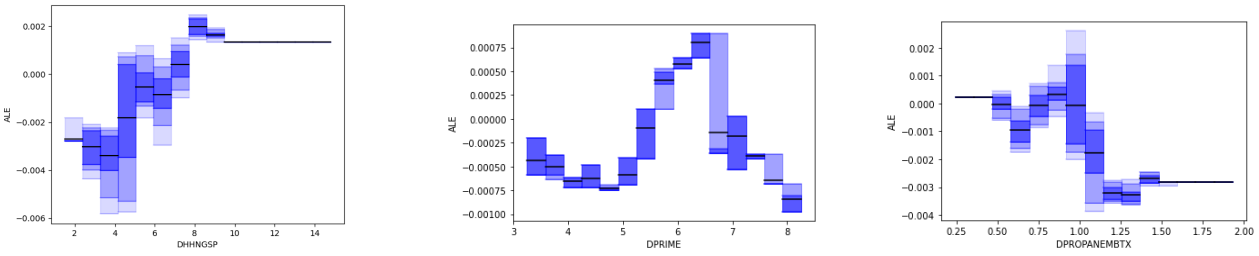

(a) DHHNGSP (b) DPRIME (c) DPROPANEMBTX

Figure 41: ALE plots of DHHNGSP, DPRIME and DPROPANEMBTX, dependent on the standard deviation of Apple

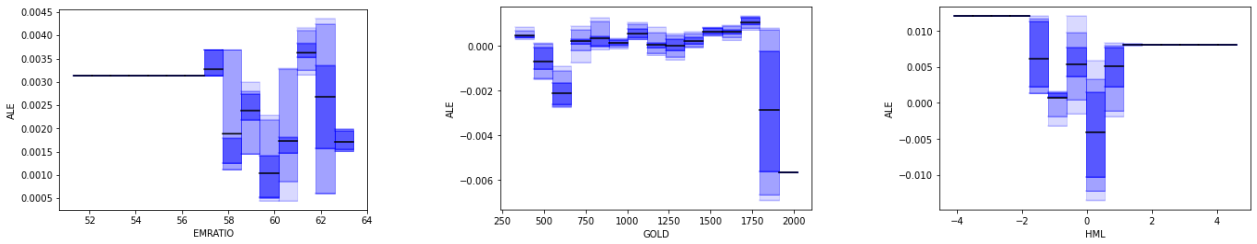

(a) EMRATIO (b) GOLD (c) HML

Figure 42: ALE plots of EMRATIO, GOLD and HML, dependent on the standard deviation of Apple

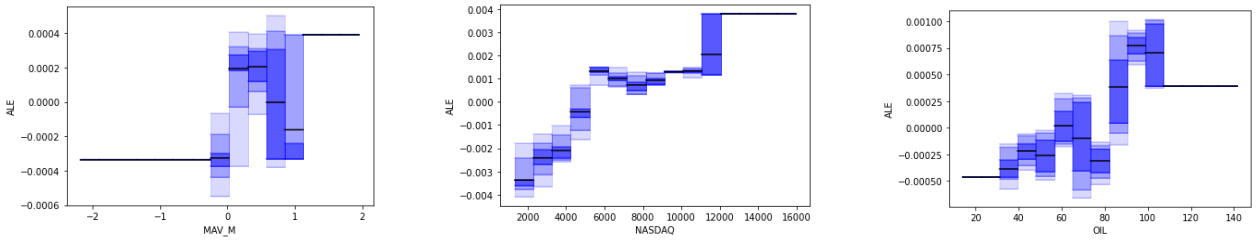

(a) MAV-M (b) NASDAQ (c) OIL

Figure 43: ALE plots of MAV-M, NASDAQ and OIL, dependent on the standard deviation of Apple

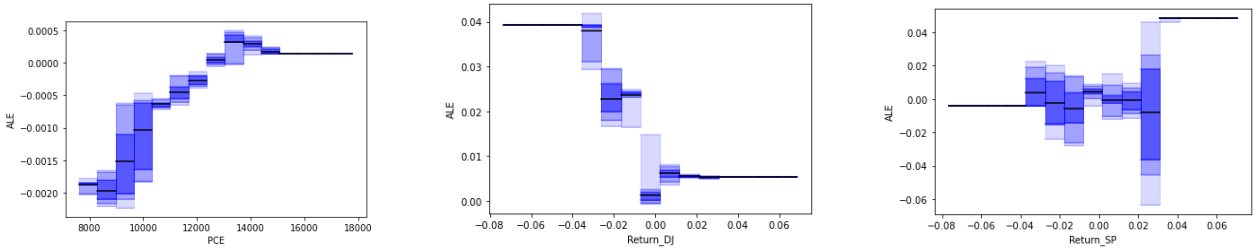

(a) PCE (b) Return-DJ (c) Return-SP

Figure 44: ALE plots of PCE, Return-DJ and Return-SP, dependent on the standard deviation of Apple

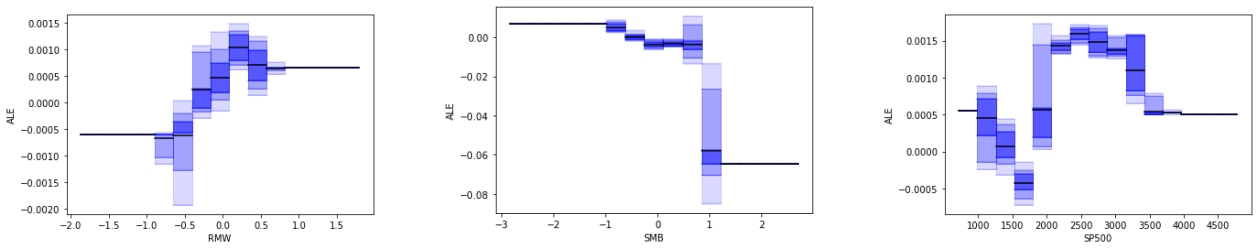

(a) RMW (b) SMB (c) SP500

Figure 45: ALE plots of RMW, SMB and SP500, dependent on the standard deviation of Apple

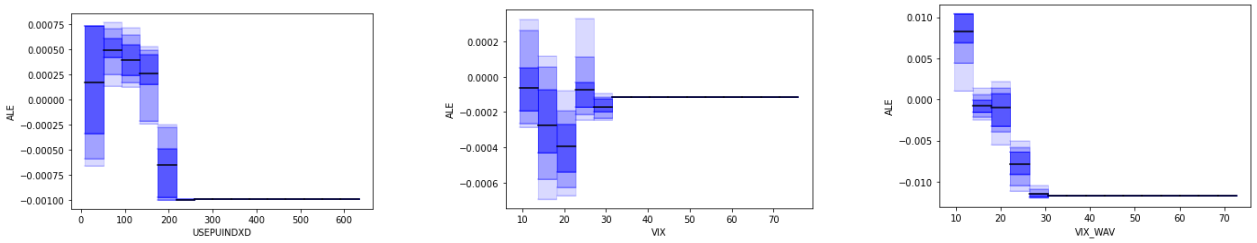

(a) USEPUINDXD (b) VIX (c) VIX-WAV

Figure 46: ALE plots of USEPUINDXD, VIX and VIX-WAV, dependent on the standard deviation of Apple

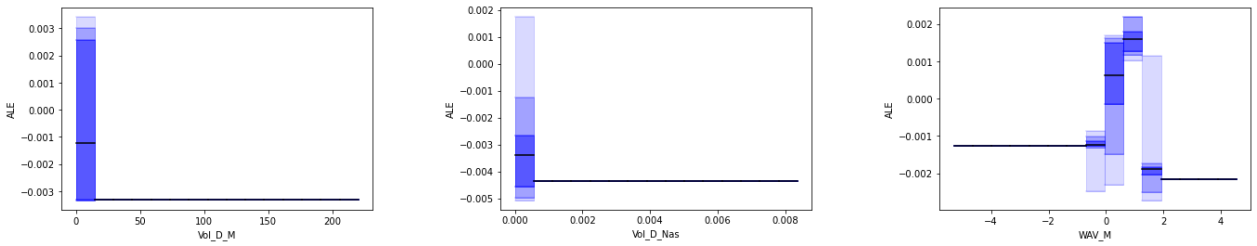

(a) Vol-D-M (b) Vol-D-Nas (c) WAV-M

Figure 47: ALE plots of Vol-D-M, Vol-D-Nas and WAV-M, dependent on the standard deviation of Apple

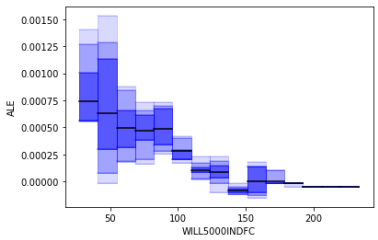

(a) WILL5000INDFC

Figure 48: ALE plots of WILL5000INDFC, dependent on the standard deviation of Apple

A.1.5 Microsoft mean

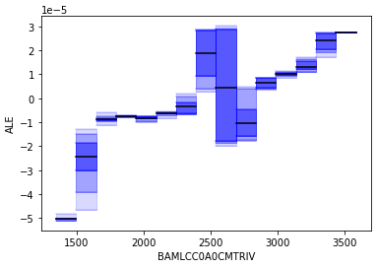

(a) BAMLCC0A0CMTRIV

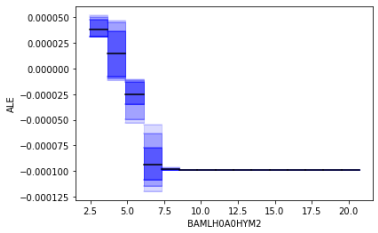

(b) BAMLH0A0HYM2

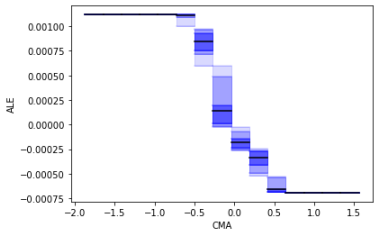

(c) CMA

Figure 49: ALE plots of BAMLCC0A0CMTRIV, BAMLH0A0HYM2 and CMA, dependent on the mean of Microsoft

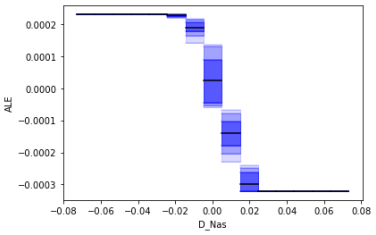

(a) D-Nas

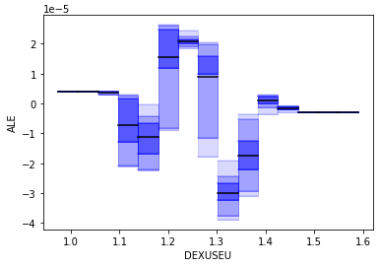

(b) DEXUSEU

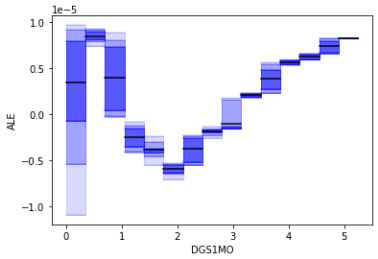

(c) DGS1MO

Figure 50: ALE plots of D-Nas, DEXUSEU and DGS1MO, dependent on the mean of Microsoft

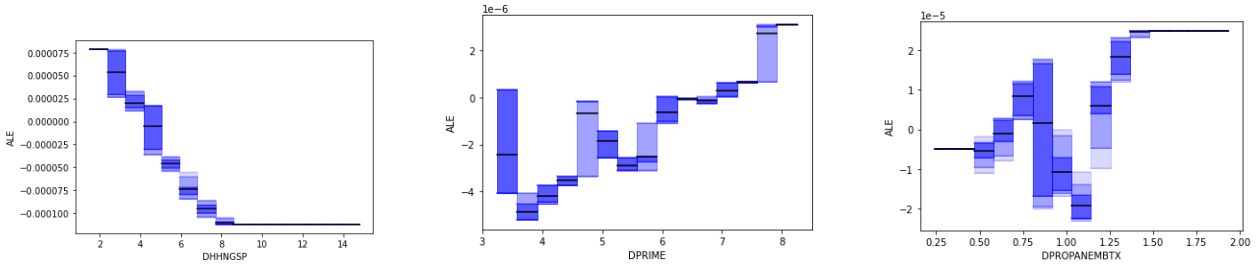

(a) DHHNGSP (b) DPRIME (c) DPROPANEMBTX

Figure 51: ALE plots of DHHNGSP, DPRIME and DPROPANEMBTX, dependent on the mean of Microsoft

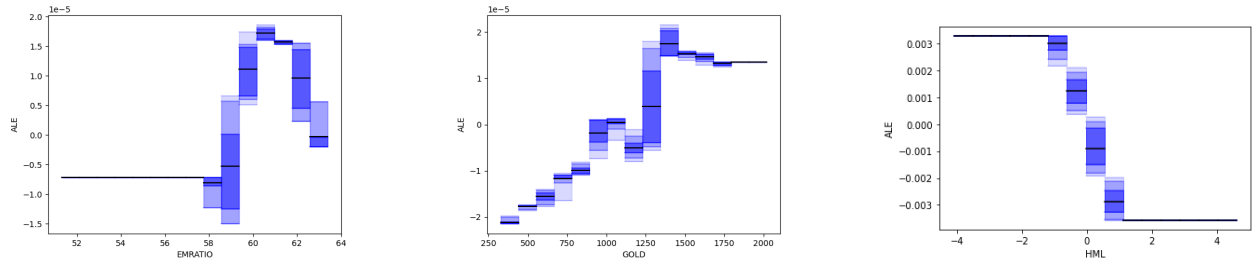

(a) EMRATIO (b) GOLD (c) HML

Figure 52: ALE plots of EMRATIO, GOLD and HML, dependent on the mean of Microsoft

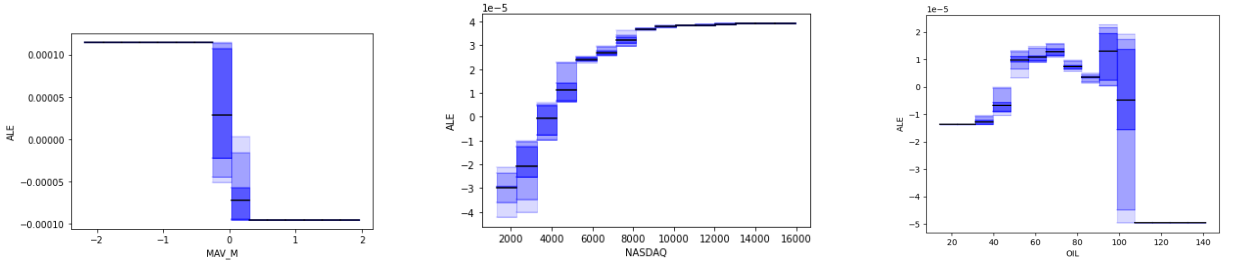

(a) MAV-M (b) NASDAQ (c) OIL

Figure 53: ALE plots of MAV-M, NASDAQ and OIL, dependent on the mean of Microsoft

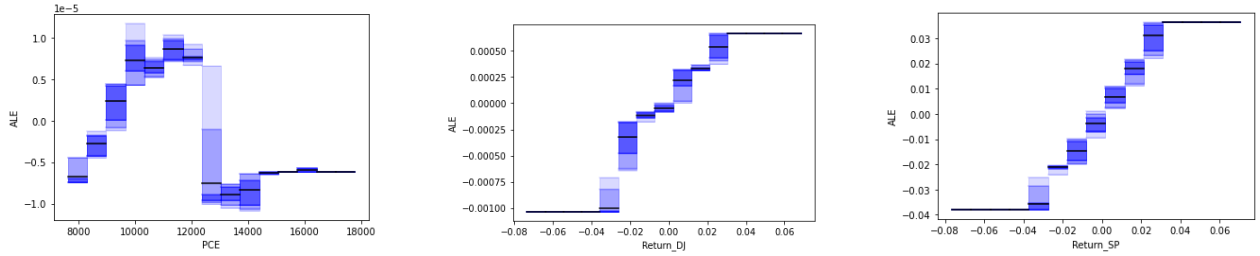

(a) PCE (b) Return-DJ (c) Return-SP

Figure 54: ALE plots of PCE, Return-DJ and Return-SP, dependent on the mean of Microsoft

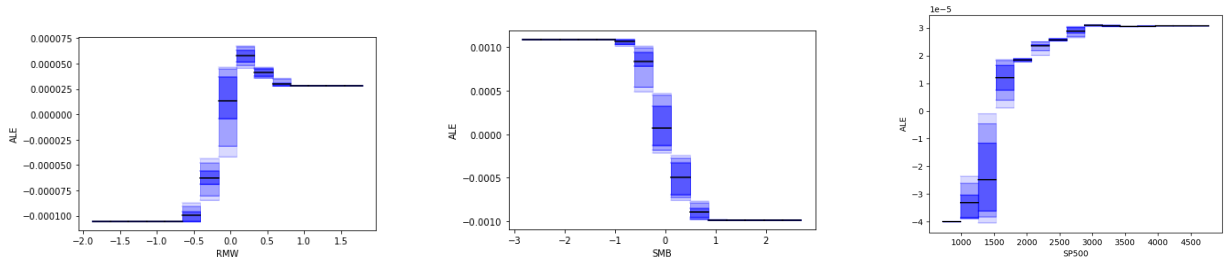

(a) RMW (b) SMB (c) SP500

Figure 55: ALE plots of RMW, SMB and SP500, dependent on the mean of Microsoft

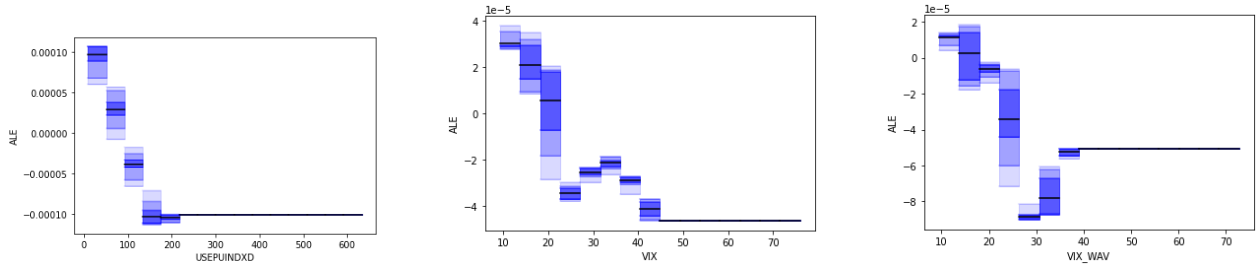

(a) USEPUINDXD (b) VIX (c) VIX-WAV

Figure 56: ALE plots of USEPUINDXD, VIX and VIX-WAV, dependent on the mean of Microsoft

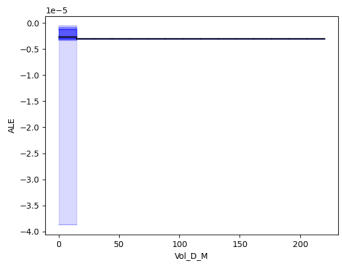

(a) Vol-D-M

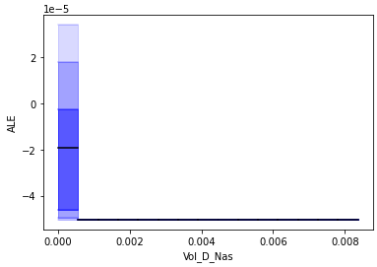

(b) Vol-D-Nas

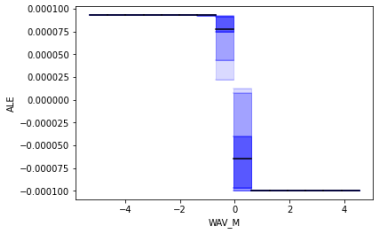

(c) WAV-M

Figure 57: ALE plots of Vol-D-M, Vol-D-Nas and WAV-M, dependent on the mean of Microsoft

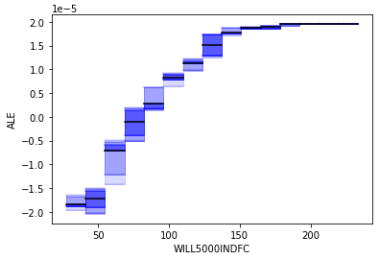

(a) WILL5000INDFC

Figure 58: ALE plots of WILL5000INDFC, dependent on the mean of Microsoft

A.1.6 Microsoft standard deviation

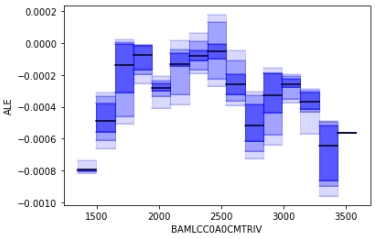

(a) BAMLCC0A0CMTRIV

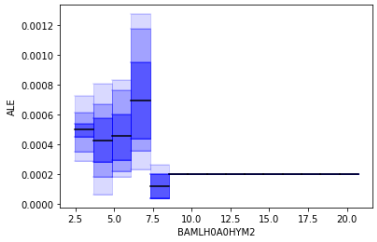

(b) BAMLH0A0HYM2

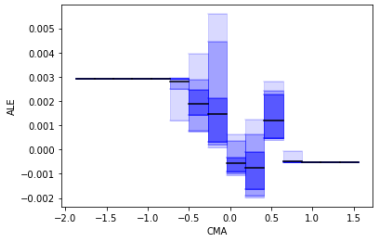

(c) CMA

Figure 59: ALE plots of BAMLCC0A0CMTRIV, BAMLH0A0HYM2 and CMA, dependent on the standard deviation of Microsoft

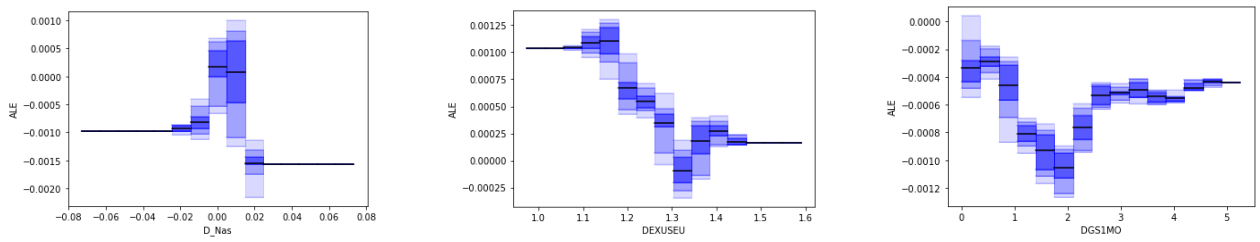

(a) D-Nas (b) DEXUSEU (c) DGS1MO

Figure 60: ALE plots of D-Nas, DEXUSEU and DGS1MO, dependent on the standard deviation of Microsoft

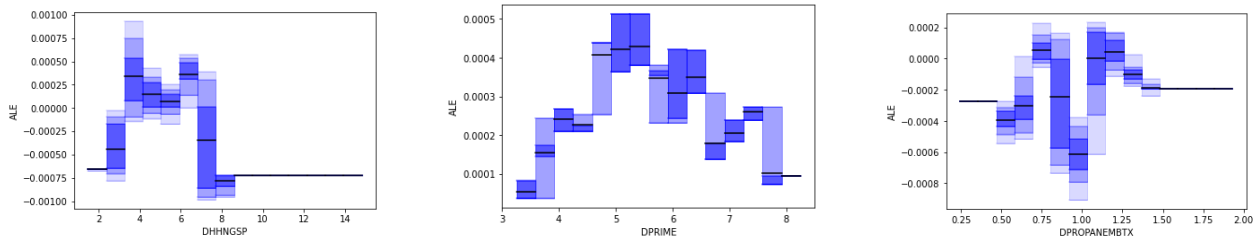

(a) DHHNGSP (b) DPRIME (c) DPROPANEMBTX

Figure 61: ALE plots of DHHNGSP, DPRIME and DPROPANEMBTX, dependent on the standard deviation of Microsoft

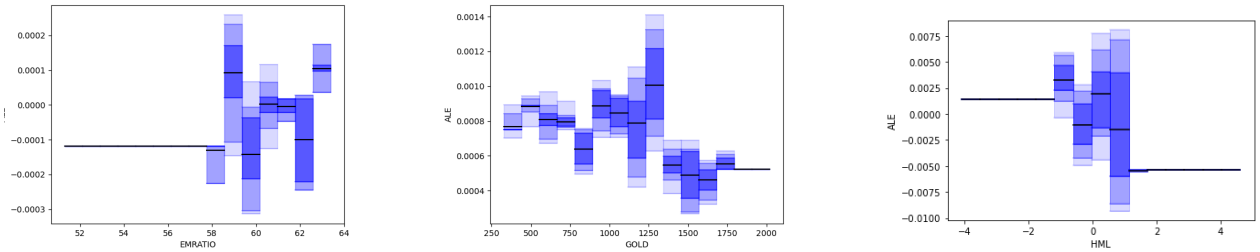

(a) EMRATIO (b) GOLD (c) HML

Figure 62: ALE plots of EMRATIO, GOLD and HML, dependent on the standard deviation of Microsoft

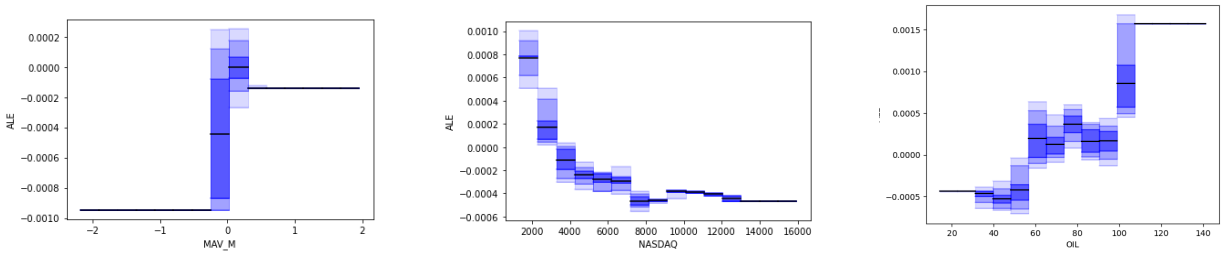

(a) MAV-M (b) NASDAQ (c) OIL

Figure 63: ALE plots of MAV-M, NASDAQ and OIL, dependent on the standard deviation of Microsoft

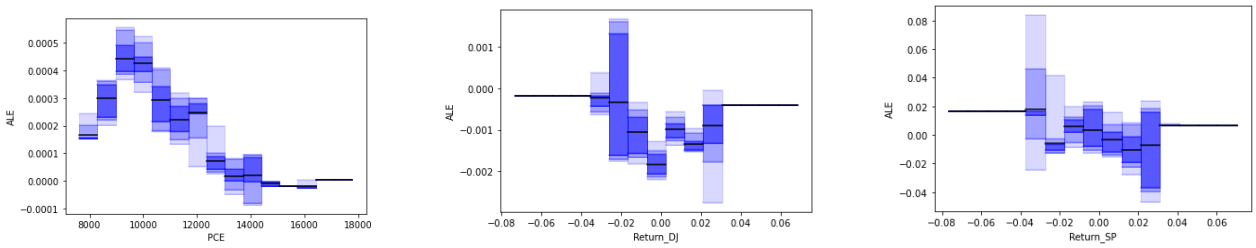

(a) PCE (b) Return-DJ (c) Return-SP

Figure 64: ALE plots of PCE, Return-DJ and Return-SP, dependent on the standard deviation of Microsoft

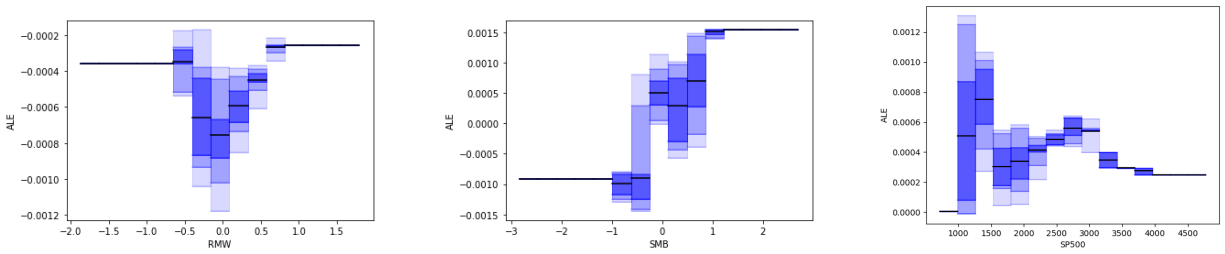

(a) RMW (b) SMB (c) SP500

Figure 65: ALE plots of RMW, SMB and SP500, dependent on the standard deviation of Microsoft

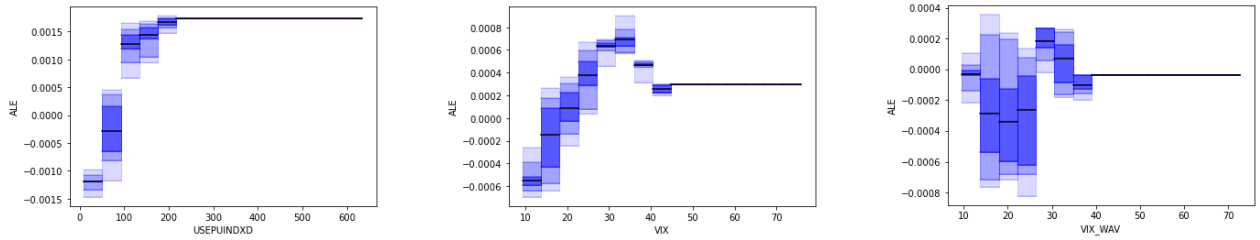

(a) USEPUINDXD (b) VIX (c) VIX-WAV

Figure 66: ALE plots of USEPUINDXD, VIX and VIX-WAV, dependent on the standard deviation of Microsoft

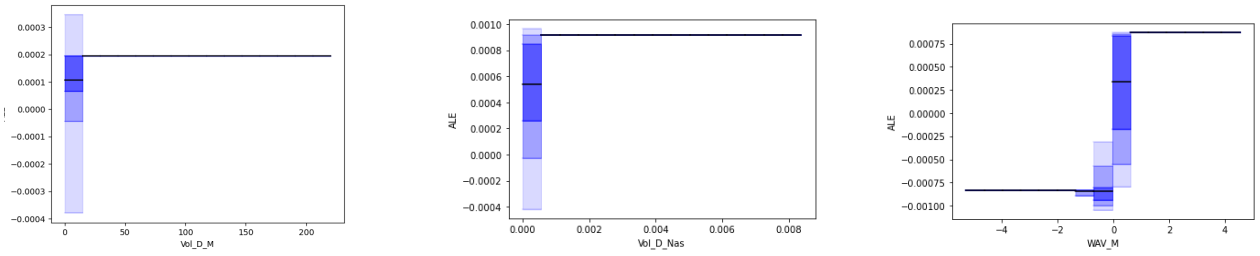

(a) Vol-D-M (b) Vol-D-Nas (c) WAV-M

Figure 67: ALE plots of Vol-D-M, Vol-D-Nas and WAV-M, dependent on the standard deviation of Microsoft

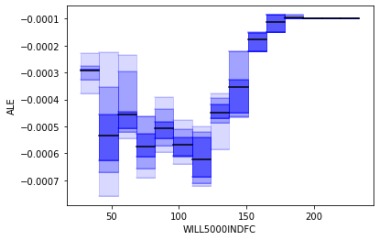

(a) WILL5000INDFC

Figure 68: ALE plots of WILL5000INDFC, dependent on the standard deviation of Microsoft