INSTITUTO
UNIVERSITÁRIO
DE LISBOA

Can we improve the accuracy of Value-at-Risk models using liquidity risk?

Henrique Manuel Gonçalves Barbeito Costa

Master in Finance

Supervisor:
PhD, António Manuel Rodrigues Guerra Barbosa, Assistant Professor, ISCTE-IUL
Iscte Business School

September, 2024

BUSINESS
SCHOOL

Department of Finance

Can we improve the accuracy of Value-at-Risk models using liquidity risk?

Henrique Manuel Gonçalves Barbeito Costa

Master in Finance

Supervisor:
PhD, António Manuel Rodrigues Guerra Barbosa, Assistant Professor, ISCTE-IUL
Iscte Business School

September, 2024

Acknowledgements

Firstly, I would like to thank my supervisor, professor António Barbosa. Without his help, navigating through this work would have proved an unsurmountable task. From aiding me in data selection, to data trimming, to challenging me on my analysis, I deeply thank his help. His input and aid have been most valuable for the completion of this thesis.

I would also like to thank my friends and family, who never stopped supporting me during this journey. A special thank you to Valériane Starte for her unconditional support and undying encouragement. Without her help I wouldn't have been able to complete this task. Thank you.

We can and will always celebrate every life event, be it good or bad, with those that are still with us. As such, this work is ultimately dedicated to my father, Manuel Costa. Even though I won't be able to share yet another step on my academic and professional journey with him, I can still feel his support to keep going.

Resumo

A métrica de medição de risco mais geralmente usada, o Value-at-Risk (VaR), subestima o risco que os investidores e as instituições enfrentam ao assumir a existência de mercados perfeitos, onde os preços dos ativos não são impactados pelas ações dos agentes de mercado. Na prática, tanto os investidores como as instituições enfrentam uma miríade de riscos que não são completamente capturados pela estimação tradicional do VaR, a qual considera apenas risco de mercado. Nesta dissertação exploramos como se conseguem produzir resultados mais precisos para ativos ilíquidos ao medir e incorporar o risco de liquidez na estimação do VaR. Deste modo, utilizamos os modelos Generalized Auto Regressive Conditional Heteroskedasticity (GARCH), student t GARCH, e o Fractionally Integrated GARCH (FIGARCH) de modelação de volatilitidade para o cálculo do VaR. Além disso, uma vez que a literatura sobre o VaR ajustado pela liquidez tende a focar-se no VaR normal paramétrico, expandimos a análise para incluir os modelos volatility-adjusted historical VaR e Monte Carlo VaR. A performance de cada modelo é avaliada através de um backtest usando os testes Unconditional Coverage e BCP. Os resultados deste estudo demonstram que ignorar o efeito da liquidez pode levar à subestimação do risco em ativos ilíquidos.

Palavras-chave: Value-at-Risk, liquidez, bid-ask spread, backtesting

Classificação JEL: C10, G32

Abstract

The industry standard risk measurement metric, the Value-at-Risk (VaR), understates the actual risk investors and institutions face by assuming perfect markets, where asset prices are unaffected by market players' actions. In practice, investors and institutions are faced with a multitude of risks which aren't fully captured with the standard VaR estimation approach, which solely captures market risk. In this dissertation we'll explore how modelling liquidity risk and incorporating it into the VaR estimation can produce more accurate measurements of risk for illiquid assets. We'll use the standard Generalized Auto Regressive Conditional Heteroskedasticity (GARCH) volatility model, as well as the student t GARCH and the Fractionally Integrated GARCH (FIGARCH) for the VaR estimation. Furthermore, since the liquidity-adjusted VaR literature tends to focus on the parametric normal VaR model, we extend the analysis to include the volatility-adjusted historical VaR and Monte Carlo VaR models. Each model's performance will be assessed though a backtest using the Unconditional Coverage and BCP test. The results from this study show that ignoring the liquidity effect can produce an underestimation of risk for illiquid assets.

Keywords: Value-at-Risk, liquidity, bid-ask spread, backtesting

JEL Classification: C10, G32

Table of Contents

Figure Index

Table Index

Introduction

With each passing year, financial institutions are faced with ever-increasing regulations that reflect the highly volatile and complex environment in which they operate. Regulators worldwide continue to strengthen their oversight and control, as even small disruptions can create financial instability (Allen et al., 2018). This instability arises from the complex and interconnected ecosystem of the markets in which these institutions operate.

Financial institutions face a myriad of risks, some undiversifiable, such as market risk, which is shared by all market makers, and some which can be mitigated through enough operational scrutiny, such as liquidity risk or operational risk. Regulatory bodies aim to strengthen the operational infrastructure of these institutions to ensure that they maintain sound risk management practices, such as imposing regulatory capital requirements, liquidity provisions requirements, stress testing requirements, among others.

This dissertation will focus on the modelling of market risk, which arises from the fluctuations in market prices of financial assets. The foremost standard metric to measure market risk is the Value-at-Risk (VaR), used by banks, investors, and other financial practitioners worldwide. It can be defined as a statistical measure that indicates the potential loss limit of an asset/portfolio that may be breached given a pre-determined confidence level over a given future time horizon (Alexander, 2009). One of the main drawbacks of the VaR is that it is estimated under the assumption of perfect markets, wherein investors have no influence over asset prices, and as such understates the actual risks institutions and investors face, such as credit risk, liquidity risk, among others.

Our main focus will be to implement the cost of the liquidity risk into the VaR. We will compute the percentual VaR for two portfolios, one denoted as the illiquid portfolio and another denoted as the liquid portfolio. Each of the portfolios is comprised of 10 stocks of U.S. equity, all equally weighted, which were chosen based on their intraday traded volume for a period of 10 years, from 2nd January 2014 to 31st December 2023. The stocks with the lowest traded volume will be part of the illiquid portfolio and those with the highest traded volume will be part of the liquid portfolio. We expect assets with low trading activity to present higher illiquidity, as there are fewer buyers and sellers that are actively participating in their market. Since the liquidity is scarcer, we expect their bid-ask spread to be larger than for actively traded assets, and as such we will be using the bid-ask spread to measure the added cost of liquidity that an investor or institution would face if they were to hold these assets. Our end goal is to understand whether including the cost of liquidity in the VaR estimation provides more accurate results for illiquid assets and reduces in a significant way the exceedances of the VaR models, where an exceedance is defined as every instance where the observed portfolio loss is lower than the VaR estimate for that same day (Alexander, 2009).

To estimate the VaR, we must first model the volatility of the returns. We will be estimating the volatility through three different models: the GARCH, the student t GARCH and the FIGARCH models.

Next, we estimate the VaR using three different models: the parametric normal VaR, the volatility-adjusted historical VaR, and the Monte Carlo VaR. The cost of liquidity is computed using the relative portfolio bid-ask spread and is then added to the VaR estimates to obtain the liquidity adjusted VaR. As such, we will be evaluating 9 different models, as each VaR model is computed using each of the three differing volatilities. Each of the 9 models is estimated twice, once as the non-adjusted VaR, and another as the liquidity adjusted VaR.

Finally, we will perform a backtest to assess each model's performance. Using the Unconditional Coverage (UC) test developed by Kupiec (1995), we will determine whether the number of exceedances of each model falls in line with the expected exceedance rate. The exceedance rate is the most useful indicator to measure a model's validity since it holds significant regulatory importance and is a direct measure of risk (Jorion, 2006). Additionally, for those models that passed the UC test, we will be computing the BCP test (Berkowitz et al., 2011) to test the autocorrelation of the exceedances, from lag 1 up to lag 5.

Results show that incoporating the cost of liquidity into the VaR substantially decreases the number of exceedances presented by the illiquid portfolio for all models, while not having as substantial an impact on the liquid portfolio models. Furthermore, some non-adjusted VaR models that were rejected by the UC test are accepted once the liquidity component is included in the VaR estimation.

This dissertation is organized as follows: Chapter 1 covers the most relevant literature; Chapter 2 presents the data and portfolio details; Chapter 3 presents all the methodology employed; and Chapter 4 reveals the result of the backtest and the model selection.

1. Literature review

1.1. Liquidity

Many consider a perfectly liquid market one in which a security can be easily traded into cash without incurring in additional costs. Kyle (1985) expanded this notion by taking into account three transactional properties of markets that impact a continuous auction equilibrium: “tightness”, the cost of liquidating a position in a short amount of time; “depth”, being able to buy/sell any given amount of a security beyond the bid-ask spread; and “resiliency”, the speed from which prices recover from a random shock. This heuristic study of liquidity was based on Black’s (1971) definition and study of liquid markets. Considering the three transactional properties of markets developed by Kyle (1985), Black’s perfectly liquid market should be infinitely tight, not infinitely deep, and resilient enough so that prices tend to their underlying value.

Risk refers to the uncertainty of future outcomes. In order to understand the future value of an asset or portfolio, modelling and hedging risk is essential. The most commonly used metric for measuring market risk is the Value-at-Risk (VaR). It is a statistical measure that indicates a loss threshold of an asset/portfolio that may be breached given a pre-defined confidence level, over a selected time horizon (Alexander, 2009). The standard VaR model assumes infinite depth and perfect markets, wherein market makers have no influence over the prices of traded assets. By failing to account for the impact of liquidating a position, the standard VaR understates actual risk by ignoring the depth, tightness, and resiliency of liquidity.

According to Bangia et al. (2008), quantifying liquidity can be approached from an endogenous or exogenous perspective. Endogenous liquidity arises from the interactions between market participants and is influenced by the size of the position held, while exogenous liquidity is the result of market characteristics and is shared by all market players, unaffected by the actions of any one agent.

Endogenous models focus on optimal executions strategies with fixed or floating intraday time horizons to solve for the liquidation problem. However, the implementation of most endogenous models is not straightforward due to their complexity and can be quite computationally intensive. Conversely, exogenous models are easier to implement, as they assume that market makers are atomistic by nature, not having any impact on prices regardless of the quantities sold. In a practical context however, the existence of adverse selection and moral hazard leads to the classic lemons problem: as progressively more quantities of an asset are sold, the market will suspect that the reason behind the selling is not inherently positive.

1.2. Endogenous models

An ad hoc measure to account for the liquidity component in the VaR is an increase of the time horizon, that is, extending the holding period of less liquid positions by assuming that a lumpsum sale at the end of the liquidation period will occur. This method was proposed and used by JPMorgan in its Riskmetrics™ model (J.P. Morgan and Reuters, 1996).

If the log prices are independent and follow an elliptical distribution, the 1-day VaR can be upscaled to any t-day horizon by multiplying it by the square root of t.

Adjusting upwards the time horizon ultimately reflects a subjective estimate of the liquidation period. Random shocks in the market can increase the illiquidity of an asset, and the time scaling modification to the VaR may underestimate potential losses. Furthermore, this method assumes that all the assets of a portfolio will be linearly liquidated over the same time period, not taking into consideration that each asset has a different level of liquidity. Lastly, this method does not consider that market makers can liquidate small portions of their portfolios on a daily basis, instead assuming that the liquidation of their entire position occurs at the end of a pre-specified period.

Al Janabi (2008) developed a practical improvement over the square root of time modification when applied to a closed-form parametric VaR. The author proposes the existence of a liquidity threshold, chosen by the trader, to measure which assets can be deemed liquid and illiquid. As such, liquidity is measured and driven by the size of the position. If the par value of an asset is n times larger than the liquidity threshold, then the unwinding of the position will be affected by market risk variations for the additional t days it will take to entirely liquidate the position.

In the case of uncertainty when choosing an appropriate time horizon to scale the VaR, the scaling factor t can be defined as a function of the trading volume of an asset such that:

$$t = \frac{TOTAL\ TRADING\ POSITION\ SIZE\ OF\ ASSET}{DAILY\ TRADING\ VOLUME\ OF\ ASSET} \tag{1}$$

Accordingly, the liquidity risk factor (overnight standard deviation of the illiquid position) is:

$$\sigma_{adj} = \sigma_t \left(\sqrt{\frac{(2t+1)(t+1)}{6t}} \right) \tag{2}$$

where σ_t is the t-day standard deviation and t is the number of liquidation days defined in Equation 1.

We can then compute the liquidity adjusted VaR as:

$$L - VaR_t = VaR_t * \sqrt{\frac{(2t+1)(t+1)}{6t}} \tag{3}$$

Al Janabi (2008) highlights several advantages of this model over the standard square root of time rule. Firstly, when $t > 1$, $L - VaR_t > VaR_t$, since illiquid positions are exposed for a longer period of time to market risk. The difference between both corresponds to the residual market risk that remains when a position cannot be fully unwound due to illiquidity (Al Janabi, 2008). Furthermore, in the case that we are computing the VaR for a single day ($t = 1$), then $L - VaR_t = VaR_t$. Finally, the position is unwound at the rate that the market conditions are optimal, to effectively quantify liquidity effects.

1.3. Exogenous models

Berkowitz (2000) proposes that when the downward sloping demand curve is used to model the liquidity effect, the effect quantities have on prices cannot be separated, ensuring that both variables cannot be assessed independently. The author approached the modelling of the exogenous liquidity risk component not as a function of time, but through the downward sloping demand curve of assets. This exogenous liquidity risk is modelled through the assumption that a firm intends to raise a certain amount of cash to meet their obligations and recurs to the sale of assets to raise that cash. The issue ultimately becomes a maximization problem, where the portfolio owner must sell a given amount of assets M_t, during a given time horizon t, by affecting as little as possible the overall portfolio value:

$$\max_{\{q_t\}} E_t \left[\sum_t^T p_t q_t \right] \; subject\ to \; \sum_t^T q_t = M_t \tag{4}$$

where p_t is the price of the asset at time t and q_t is the quantity of the asset sold at time t. Bertsimas and Lo (1998) show that the optimal solution to the problem of maximization is to sell q_t^*, which is obtained through:

$$q_t^* = \frac{M}{t} \tag{5}$$

After orderly liquidation, the asset's price will depress due to, on the one hand, market-wide factors shared by all market makers (x_t), and on the other hand, possible shocks resulting from the market's reaction to the firm's actions ($-\theta q_t^*$). Hence, in equilibrium:

$$\hat{p}_{t+1} = p_t + x_{t+1} - \theta q_t^* \tag{6}$$

The objective of the portfolio manager thus becomes to forecast the portfolio's value after the sale of the assets occurs:

$$\hat{y}_{t+1} = Q_t' p_{t+1}' = Q_t'(p_t' + x_{t+1} - \theta q'^*_t) \tag{7}$$

where Q_t' is an $N \times 1$ vector of asset positions, p_t' is the $N \times 1$ vector of asset prices at equilibrium, and q'^*_t is the $N \times 1$ vector of asset quantities that must be sold to obtain profit maximization. Furthermore, $Q_t'(p_t + x_{t+1})$ represents the market risk component of portfolio risk, accounting for changes in asset values due to market movements shared by all market participants, and $Q_t'(-\theta q_t^*)$ measures the possible price drops from the market reaction to the asset sales (Berkowitz, 2000).

Through a simple linear regression, it's possible to obtain the values of the liquidity component θq_t and its variance, the latter being incorporated into the VaR model.

Bangia et al. (2008) proposed modelling order processing costs to measure the exogenous liquidity risk faced by all market makers. The authors identified two distinct types of liquidity risk: exogenous and endogenous illiquidity. Endogenous illiquidity can be controlled by market makers and is defined as being the result of the sudden unloading of large positions that the market cannot easily absorb. This effect is driven by the size of the position – the larger the size, the larger the illiquidity, as transforming such a large position into cash would prove progressively more difficult. On the other hand, exogenous illiquidity is common to all market participants and is undiversifiable. This exogenous liquidity risk can be observed through the bid-ask spread. The costs arising from market transactions ultimately dictate how much a trader is willing to buy or sell a given asset, thus influencing the bid and ask prices in the market.

As such, investors are faced with two types of risk: market risk, which is undiversifiable; and liquidity risk, which can in turn be separated into exogenous and endogenous illiquidity. Bangia et al. (2008) measure exogenous illiquidity and incorporate it into the VaR model through the bid-ask spread.

When trying to liquidate a position quickly, or in adverse market conditions, traders don't realize the mid-price when closing their order, they realize the mid-price minus the bid-ask spread (Bangia et al., 2008). Under this condition, mark to market pricing understates the actual risks faced as it assumes that all positions will be closed at the mid-price. According to Bangia et al. (2008), the key to implementing exogenous illiquidity into the standard VaR is to model the distribution of the deviation of the actual closing price from the mid-price.

If the size of the position to be liquidated is lower than the quote depth, then a transaction is closed at the quote price, with a cost of immediate execution of half of the bid-ask spread.

Since the market has enough open orders to match the position that is going to be sold, the transaction will only be affected by exogenous risk.

However, if the size of the position to be sold is higher than the quote depth, then the cost of the trade will be higher than the observable bid-ask spread. Furthermore, since endogenous risk is driven by the size of a position, this trade will be impacted by both exogenous and endogenous illiquidity.

Figure 1.1 shows the relationship between the total position size and the liquidation price. As the position size increases, the liquidation will be increasingly affected by endogenous illiquidity, which can be identified by a larger bid-ask spread, and the cost of the trade will be higher than the observable bid-ask spread.

Figure 1.1

Effect of position size on liquidation value

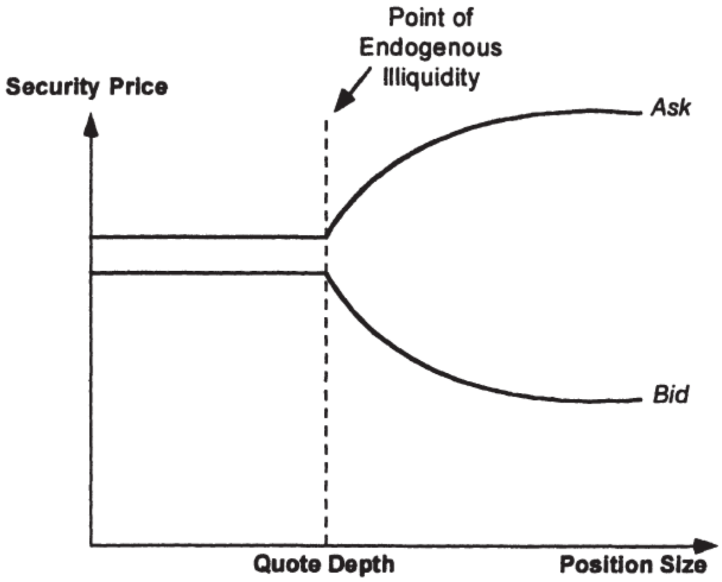

Note. Retrieved from Bangia et al. (2008). *Modeling Liquidity Risk with Implications for Traditional Market Risk Measurement and Management.* NYU Working Paper No. FIN-99-062.

To illustrate the computation of the VaR, we will be focusing on the standard parametric normal VaR model, which is further described in Section 3.2. alongside the other VaR models that are used in our case study. As such, the $100\alpha\%$ parametric normal VaR is defined as:

$$VaR_{t,\alpha} = -\Phi^{-1}(\alpha) \times \sigma_t \tag{8}$$

where $\Phi^{-1}(\alpha)$ is the standard normal α-quantile value and σ_t is the second conditional moment of the mid-price log returns at time t.

The above parametric model only considers the distribution of the mid-price. On average, it is expected for the bid price to be less than the mid-price by half the average relative spread, $\frac{1}{2} * \bar{S}$ (Bangia et al., 2008). Hence, since liquidation is not realized at the mid-price, the average parametric model needs to be modified to consider the exogenous cost of liquidity, which explains the deviation from the mid-price when settling a position.

Formally, the exogenous cost of liquidity (COL_t) is defined as:

$$COL_t = \frac{1}{2}[P_t(\bar{S} + a * \tilde{\sigma}_t)] \tag{9}$$

where P_t is the observable mid-price at time t, $\bar{S}$ is the average bid-ask relative spread, defined as $([Ask_t - Bid_t]/Mid_t)$, $\tilde{\sigma}_t$ is the volatility of the relative spread at time t, and a is a scaling factor that ensures $(1 - \alpha)\%$ probability coverage, where α corresponds to the significance level of the VaR model being estimated.

When examining the statistical properties of the bid-ask spread, Affleck-Greaves et al. (2000) conclude that the raw spread, relative spread, and changes in spread present large deviations from normality. As such, because the cost of liquidity is based on the average relative spread, which is commonly far from normal, the scaling factor a cannot be inferred based on a multivariate Gaussian distribution.

Bangia et al. (2008) make the simplifying assumption that extreme spread events and extreme return events occur simultaneously. This assumption is sustained by the fact that risk managers are not interested in analysing and modelling average spreads, but rather extreme spreads in adverse market conditions. Furthermore, in periods of increased market volatility, the correlations between extreme spreads and extreme returns can be shown to be high enough to warrant a strong observable relationship between return risk and exogenous liquidity risk.

The parametric normal liquidity adjusted VaR ($L - VaR$) for a single asset is thus:

$$L - VaR_t = VaR_{t,\alpha} + COL_t = -\Phi^{-1}(\alpha) \times \sigma_t + \frac{1}{2}[P_t(\bar{S} + a * \tilde{\sigma})] \tag{10}$$

The traditional portfolio level VaR uses the covariance matrix of returns to model portfolio risk. However, using the covariance matrix for the $L - VaR$ would imply that the spreads follow a normal multivariate distribution, which may not accurately represent their actual distribution. The approach that we will follow is to compute the weighted average bid, ask, and mid prices at portfolio level, and use them to compute the relative average spread and the portfolio level cost of liquidity.

2. Data and descriptive statistics

This study aims to examine the impact of modelling the liquidity effect of assets and incorporating it into the VaR. The object of study will be 2 different portfolios, each comprised of long positions of United States (U.S.) equity. One of the portfolios will be composed of highly liquid stocks, and the other portfolio will be composed of less liquid stocks, both based on the intraday average traded volume for a period of 10 years, from 2nd January 2014 to 31st December 2023. The local currency used will be the U.S. Dollar (USD).

Each portfolio is composed of 10 stocks that are listed on the S&P500, all equally weighted, with each stock representing 10% of the total portfolio weight. The liquidity component will be modelled based on the bid-ask spread of the shares, which was obtained from Bloomberg. Timewise, we will be using daily data for a period of 10 years, from 2nd January 2014 to 31st December 2023. Table 2.1 shows the portfolio composition for each of the portfolios.

Table 2.1

Portfolio composition

Stock	Ticker	Weight
Teledyne Technologies	TDY	10%
Tyler Technologies Inc	TYL	10%
Zebra Technologies	ZBRA	10%
Jack Henry & Associates, Inc.	JKHY	10%
Bio-Rad Laboratories, Inc.	BIO	10%
Teleflex	TFX	10%
News Corp	NWS	10%
Nordson Corporation	NDSN	10%
Packaging Corp Of America	PKG	10%
Marketaxess Holdings Inc	MKTX	10%
Illiquid Portfolio Total		**100%**
Apple	AAPL	10%
Advanced Micro Devices	AMD	10%
Nvidia	NVDA	10%
Bank of America Corp	BAC	10%
Wells Fargo & Co	WFC	10%
Pfizer	PFE	10%
Bristol-Myers Squibb Co	BMY	10%
Freeport-McMoRan	FCX	10%
AT&T Inc	T	10%
American Airlines Group	AAL	10%
Liquid Portfolio Total		**100%**

Note. This table shows each stock and its respective weight in both the liquid and illiquid portfolio.

According to Bangia et al. (2008), in order to expand the liquidity adjusted VaR from a single financial instrument level to a portfolio level, we would need to use the covariance matrix of the assets' returns and the bid-ask spread covariance matrix, the latter needed to quantify the liquidity effect. However, this imposes the assumption that spreads follow a multivariate normal distribution, which is rendered unrealistic as the spread distribution isn't as well behaved as the return distribution (Affleck-Greaves et al., 2000). The solution proposed by the authors is to calculate the portfolio level bid and ask prices as the weighted average of the individual instruments. Hence, as we retrieve the individual bid and ask prices of each stock, we compute the weighted average bid and ask prices at portfolio level for each trading day. The portfolio mid prices, which are extrapolated as the average of the bid and ask quotes at the close of trading day t, are used to compute the daily log returns, which allow us to estimate the daily VaR in percentage points:

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right) \tag{11}$$

where r_t represents the return at time t and P_t is the portfolio mid-price at time t.

Table 2.2 displays the descriptive statistics of each portfolio's daily returns. As is expected from financial return data (Fama, 1965), the returns for each portfolio show excess kurtosis (above 3), and negative asymmetry. As such, the unconditional distribution of the returns seems to deviate from a multivariate normal distribution, and using other non-normal distributions to fit the data may be more appropriate.

Table 2.2

Descriptive statistics of each portfolio's returns

Returns	Mean	Median	Max	Std Deviation	Skewness	Kurtosis
Illiquid Portfolio	0.000498	0.000909	0.091059	0.012510	-0.321822	5.609848
Liquid Portfolio	0.000103	0.001080	0.117231	0.027402	-20.189947	615.131034

Note. Mean, median, maximum, standard deviation, skewness, and kurtosis of the daily returns for each portfolio, from 2nd January 2014 to 31st December 2023.

It is also useful to analyze the distribution of the portfolio spreads, as we expect them to deviate from normality. Table 2.3 displays the descriptive statistics of the relative spreads of each portfolio. Both portfolios' percentage spreads present very large kurtosis and positive asymmetry, suggesting that both follow a non-normal distribution. As expected of less liquid stocks, the illiquid portfolio exhibits a larger average relative spread when compared to the liquid portfolio, accounting for the added costs of holding these assets.

Table 2.3

Descriptive statistics of each portfolio's relative spreads

Relative Spread	Mean	Median	Max	Std Deviation	Skewness	Kurtosis
Illiquid Portfolio	0.052313	0.042589	1.723732	0.047270	18.637331	623.455412
Liquid Portfolio	0.026203	0.024699	0.113590	0.011892	0.881406	24.171353

Note. Mean, median, maximum, standard deviation, skewness, and kurtosis of the daily relative spreads for each portfolio, from 2nd January 2014 to 31st December 2023.

To verify that the returns do not follow a normal distribution, the Jarque-Bera and the Shapiro-Wilk tests for normality were performed.

The Jarque-Bera test (Jarque & Bera, 1987) is a goodness-of-fit test that tests the joint null hypothesis of the skewness and the excess kurtosis of the observations being 0. The Jarque-Bera test makes use of the Lagrange Multiplier test to check the normality of the observations, such that the test statistic is given by:

$$LM = n\left(\frac{S^2}{6} + \frac{(K-3)^2}{24}\right) \tag{12}$$

where n is the number of observations in the sample, S is the sample skewness and K is the sample kurtosis. This test statistic is asymptotically distributed as a $\chi^2_{(2)}$ under normality, as it is the sum of squares of two asymptotically independent standard normal distributions (Bowman & Shenton, 1975).

The Shapiro-Wilk test (Shapiro & Wilk, 1965) tests the null hypothesis that the observations in the sample come from a normally distributed population. The test statistic is given by:

$$W = \frac{\left(\sum_{i=1}^{n} a_i X_{(i)}\right)^2}{\sum_{i=1}^{n}(X_i - \bar{X})^2} \tag{13}$$

where $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(n)}$ are the ordered values of the vector of random variables $X = (X_1, \dots, X_n)$, $\bar{X}$ is the mean of the sample, and a_i are tabulated coefficients.

The results of the tests are displayed in Table 2.4. As is expected, for the returns of each portfolio we fail to accept the null hypothesis of both the Jarque-Bera and the Shapiro-Wilk test and can thus expect the returns to follow a non-normal distribution.

Table 2.4

Jarque-Bera and Shapiro-Wilk test results

	Test	*p*-value	Result
Illiquid portfolio	Jarque-Bera	0.00000**	Non-Normal
	Shapiro-Wilk	0.00000**	Non-Normal
Liquid portfolio	Jarque-Bera	0.00000**	Non-Normal
	Shapiro-Wilk	0.00000**	Non-Normal

Note. Allows to conclude about the normality of the return distribution. ** denotes that the null hypothesis was rejected at the 1% significance level.

3. Methodology

3.1. Volatility models

In the context of financial data analysis, volatility (σ) represents the degree of variation in the prices of financial instruments over time. In this thesis we are interested in modelling the conditional volatility of daily returns, σ_t. To model the volatility, we use three different models: the GARCH model with normal and student t-distributed innovations, and the FIGARCH model.

For each model, we compute daily volatility estimates using the first 252 observations (i.e., 1 trading year) as training data. Furthermore, to ensure that the estimates reflect the current market conditions, we re-estimate the model parameters every 21 trading days.

3.1.1. GARCH model

Bollerslev (1986) introduced a useful generalization of the standard Autoregressive Conditional Heteroskedasticity (ARCH) model proposed by Engle (1982). Both models allow the conditional volatility to change over time as a function of past errors while the unconditional volatility remains unchanged.

The GARCH model captures the effects of long-term memory and volatility clustering more effectively than the standard ARCH model by incorporating both the lagged squared residuals and the lagged conditional variances. By including the lagged conditional variances as a variable, the GARCH model exhibits an adaptive learning behaviour. Furthermore, the GARCH model exhibits excess kurtosis which makes it useful for modelling the conditional volatility and heavy tails of financial market data simultaneously.

The series of daily returns, r_t, is modelled as follows:

$$r_t = \mu + \varepsilon_t,\ \varepsilon_t \sim N(0, \sigma_t^2) \tag{14}$$

where μ is the mean value of the returns and ε_t is the residual at time t. This error term represents the deviation of the observed return r_t from the mean μ.

The GARCH (p,q) model defines the variance as:

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2 \tag{15}$$

where ω is a constant term in the regression, α_i measures the impact of past squared errors, and β_j measures the persistence of past variances. For the model to be valid, the parameters must be subject to the following conditions: $\omega > 0$, $\alpha_i \geq 0$ and $\beta_j \geq 0$.

Since the only data available is r_t, the series of daily returns, all the GARCH model parameters are computed simultaneously using maximum likelihood. The p and q variables represent the order of the autoregressive component and the order of the moving average component, respectively. Thus, p captures the lagged conditional variances while q captures the lagged squared residuals.

Even though several lag orders can be defined, Hansen and Lunde (2005) demonstrate that the GARCH (1,1) model isn't outperformed by other volatility models of higher order. Since the GARCH (1,1) model is less computationally expensive and demonstrates comparable performance to higher order GARCH models, it will be used in this dissertation. It will also be used for every other volatility model employed in this dissertation, as they are all extensions of the GARCH model.

Even though the leptokurtic nature of the GARCH model may be useful for modelling heavy tails in financial data series, the assumption of normality in the error term limits the model from properly modelling fat tails in the distribution. Bollerslev (1987) proposes a simple extension of the ARCH and GARCH models to allow the errors to be conditionally t-distributed.

The Student's t-distribution is symmetric around 0 and has fatter tails than the Gaussian distribution, making it more appealing for the study of financial data (Bradley & Taqqu, 2003). In fact, several studies such as Fielitz and Roselle's (1983) and Boothe and Glassman's (1987) indicate that non-normal distributions may be preferred when dealing with financial datasets, as they accommodate more easily skewed distributions and other higher distribution moments.

The t-distributed GARCH (p,q) model for the conditional variance is defined as:

$$\sigma_{t|t-1}^2 = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j|t-1-j}^2 \tag{16}$$

where $\varepsilon_t|\psi_{t-1} \sim f_v(\varepsilon_t|\psi_{t-1})$.

As with the standard GARCH model, we will be using the GARCH (1,1)-t model.

3.1.2. FIGARCH model

The Integrated GARCH (IGARCH) was developed by Engle and Bollerslev (1986) as a solution to the short-term memory of standard linear models. As a non-linear model, the autocorrelation function of the IGARCH decays slowly over time, having infinite memory. However, this means that any shock that occurs within the sample will never die out, making the estimators biased.

As an improvement to the infinite persistence present in the IGARCH, Baillie et al. (1996) developed the Fractionally Integrated GARCH (FIGARCH) as a long memory model that allows a slow hyperbolic rate of decay for the lagged squared or absolute innovations. This is achieved by incorporating the fractional differencing parameter d into the IGARCH model. This parameter allows the model to handle persistence in the volatility more flexibly, allowing it to decay slowly over time. For the model to be valid, it is necessary that $0 \leq d \leq 1$.

The FIGARCH (p,d,q) conditional volatility is thus:

$$\sigma_t^2 = (\omega - \bar{\varepsilon}_t^2) + \sum_{i=1}^{q} \alpha_i(\varepsilon_{t-i}^2 - \bar{\varepsilon}_{t-i}^2) + \sum_{j=1}^{p} \beta_j\left(\sigma_{t-j}^2 - \bar{\varepsilon}_{t-j}^2\right) \tag{17}$$

where ε_t^2 is the squared error, and $\bar{\varepsilon}_t^2$ is the squared average error.

The model we will be using is the FIGARCH (1,d,1) model.

3.2. Value-at-Risk Models

The Value-at-Risk (VaR) can be defined as the maximum expected loss of an asset/portfolio over a selected time horizon h, for a given confidence level α. Moving forward we will be using a significance level of $\alpha = 5\%$ (which corresponds to a confidence level $1 - \alpha = 95\%$), and a time horizon of 1 day ($h = 1$). As such, when applied to our data, we are $(1 - \alpha)\%$ confident that the observed losses won't exceed the VaR estimates in the next day. For more frequently traded assets such as stocks, the 1-day time horizon can accurately reflect their liquidity, as it assures that the risk assessments can be reflective of actual market conditions (Alexander, 2009). It also reduces the need for market assumptions over longer periods, which can induce potential errors in risk estimates.

In mathematical terms, the $100\alpha\%$ h-day Value-at-Risk ($VaR_{h,\alpha}$) is minus the α-quantile of the h-day return distribution. As mentioned by Alexander (2009), the α-quantile of the h-day distribution of a random variable X, denoted by $x_{h,\alpha}$, is defined by:

$$P\left(X_h < x_{h,\alpha}\right) = \alpha \tag{18}$$

Given the probability function of X is known, $x_{h,\alpha}$ can be obtained as:

$$x_{h,\alpha} = F^{-1}(\alpha) \tag{19}$$

where F^{-1} is the inverse cumulative distribution function of X.

As we expect our maximum loss to be exceeded for a given probability of α, the VaR is thus:

$$VaR_{h,\alpha} = -x_{h,\alpha} \tag{20}$$

Assuming a stationary portfolio composition for the entire holding period, which starts on 2nd January 2014 and ends on 31st December 2023, our goal is to generate a historical series of daily VaR estimates for each model in percentage terms. In this dissertation, the estimation methods that will be used are the parametric normal VaR, volatility-adjusted historical VaR and Monte Carlo simulation VaR. Each model will be estimated three times, alternating between the three volatility methods explained in Section 3.1. The liquidity adjusted VaR will be obtained by adding the exogenous cost of liquidity to each of the 9 VaR models. The series of daily VaR estimates for each model will be calculated for the 8-year period starting on January 5, 2016 and ending in December 29, 2023.

3.2.1. Normally distributed parametric VaR

Parametric VaR methods assume a parametric distribution of the portfolio returns. Even though these methods are less computationally expensive, as they rely solely on computing the distribution parameters, they impose a parametric distribution over the portfolio returns which may not appropriately model their behavior.

The parametric normal VaR makes the simplifying assumption that the portfolio returns, denoted as a continuous random variable X, are i.i.d. and follow a normal distribution. Thus, we assume $X_h \overset{i.i.d.}{\sim} N(\mu_h, \sigma^2{}_h)$, where μ and σ^2 are forecasts of the mean and variance of the returns, respectively.

Based on Equation 20, and noting that the returns follow a normal distribution, it gives that:

$$VaR_{h,\alpha} = -\Phi^{-1}(\alpha) \times \sigma_h - \mu_h \tag{21}$$

where $-\Phi^{-1}(\alpha)$ is the standard normal α-quantile value.

According to Kim et al. (1999), the drift μ_h can be dropped for shorter horizons as mean forecasts aren't likely to produce accurate predictions of future returns. Furthermore, since the volatility σ_h can be shown to produce higher estimates than the expected returns over short horizons, the forecasted returns' distribution is drowned out by the volatility estimates on the short term, indicating that μ_h won't have a significant effect on the VaR. As such, since we are dealing with daily return data ($h = 1$), we will be assuming a zero-mean estimate for the portfolio returns.

By dropping the drift adjustment from Equation 21, the $100\alpha\%$ h-day VaR is:

$$VaR_{h,\alpha} = -\Phi^{-1}(\alpha) \times \sigma_h \tag{22}$$

3.2.2. Volatility-Adjusted Historical VaR

Contrary to parametric models, which assume that the returns follow a specific parametric distribution, the historical simulation method is based on the empirical distribution of the returns, using the α quantile of the empirical distribution to compute the VaR. As such, the historical simulation relies on the dependencies of the risk factors on expected returns and co-movements between risk factors, and only makes the assumption that the future return distribution will be equal to the past return distribution.

Formally, the $100\alpha\%$ h-day historical VaR is the α quantile of the empirical h-day distribution of the returns. The historical VaR is estimated by following these steps: choose the sample size, n; compute the h-day returns of the portfolio over the sample period; compute the empirical h-day returns distribution by keeping the portfolio weights constant; sort the returns in an ascending order; and, finally, identify the α quantile of interest.

This method relies on the history of past changes of risk factors over the sample size, assuming that the current portfolio was held in the past. Thus, the choice of the sample size can significantly influence the accuracy of the risk estimates. By increasing the sample size, we can improve the precision of the empirical distribution, albeit at the cost of the historical distribution not reflecting current market conditions. This issue arises from the fact that the simple historical VaR assigns the same weight to each observation in the sample. If each return is given equal weight within the sample, we assume that all observations hold the same impact over the current market conditions. This approach can lead to ghost features in the return distribution, as an extreme observation holds the same influence on the day it occurs until the day it exits the sample, even though it may no longer be relevant.

To face this issue, Hull and White (1998) propose adjusting the volatility of the entire series of returns while still assigning the same weight to each observation, so that the current volatility, and consequently the current market conditions, are properly reflected.

By adjusting the series of returns r_t using the volatility estimates ($\hat{\sigma}_t$) and today's volatility ($\hat{\sigma}_T$), we obtain:

$$\hat{r}_t = \frac{\hat{\sigma}_T}{\hat{\sigma}_t} r_t \tag{23}$$

where $\hat{r}_t$ is the adjusted return at time t, $\hat{\sigma}_t$ is the volatility estimate at time t (computed at the end of $t-1$), and T is the VaR measurement date ($t < T$).

To adjust r_t, we will need to wait until the end of day t so that the portfolio return is observed. However, at the end of day t, we already know $\hat{\sigma}_{t+1}$, which is the most current estimate of the volatility. Therefore, the quality of the volatility adjustment can be improved by using $\hat{\sigma}_{T+1}$ and $\hat{\sigma}_{t+1}$:

$$\hat{r}_t = \frac{\hat{\sigma}_{T+1}}{\hat{\sigma}_{t+1}} r_t \tag{24}$$

We define the model that uses these adjusted returns as the volatility-adjusted historical VaR. To compute the h-day volatility-adjusted historical VaR, we choose a sample size of one trading year (252 observations). This sample size is large enough to ensure that the empirical distribution of the returns can be properly estimated, whilst also considering enough observations to ensure that current market events are being reflected.

3.2.3. Monte Carlo Simulation VaR

The Monte Carlo simulation is a mathematical technique used to solve problems that may be deterministic by nature using repeated random sampling to obtain numerical results. In the context of VaR estimation, the Monte Carlo method is a non-parametric method which provides wide flexibility without assuming the distribution of the risk factors. It involves simulating the future price paths of assets a specific number of times in order to improve the modelling of uncertainty in the time series, and to provide a more in-depth explanation of the dynamic characteristics of the assets (Alexander, 2009). By exploring an extensive number of simulations, most of the possible future outcomes of the asset can be predicted, allowing the distribution of the assets at future dates to be estimated more accurately.

Computing the Value-at-Risk through a Monte Carlo simulation can be done by taking a multivariate or a univariate approach, depending on the relationship between the variables. A multivariate distribution involves modelling the joint probability distribution of two or more uniform random variables, which is useful when the data is more complex and presents more risk factors, while the univariate distribution focuses on a single variable (Alexander, 2009).

Simulating a univariate distribution consists in modelling the behavior of one single random variable, which is described by a single cumulative distribution function. Let u be a uniform random variable obtained from the interval $(0,1)$, generated from random sampling, and x be another uniform random variable with a continuous distribution function F. The value of the simulation for x is obtained by:

$$x = F^{-1}(u) \tag{25}$$

where F^{-1} is the inverse of the continuous distribution function of x. Hence, given a uniform random variable u, the corresponding simulation of x is the u quantile of its distribution.

Since we are analyzing the portfolio returns, which can be effectively treated as a univariate time series, we will rely on the univariate distribution to describe the variability of the returns.

The Monte Carlo VaR is estimated by following these steps: simulate a distribution of the portfolio's h-day returns; sort the returns in an ascending order; and identify the α quantile from the distribution of simulated returns. Therefore, the Monte Carlo VaR uses the empirical distribution of the simulated returns to obtain the $100\alpha\%$ h-day VaR.

Choosing an appropriate number of simulations is paramount to obtaining accurate estimations, reducing statistical errors, and achieving better convergence of results. However, a trade-off occurs as we increase the number of repetitions, whereby the estimation process becomes progressively more computationally expensive to factor all the possible future path scenarios of the data. For the Monte Carlo simulation we will be running, we have decided to use 50,000 simulations to ensure accurate estimations. Furthermore, as with the volatility-adjusted historical VaR, we will be using a sample size of 1 trading year for the Monte Carlo VaR estimation.

3.2.4. Liquidity Adjusted VaR

The purpose of the liquidity adjusted VaR ($L - VaR$) is to incorporate the effects of liquidity into the VaR model to fully explain total risk. As previously defined in Equation 10, $L - VaR$ is computed as the sum of the unadjusted VaR and the cost of liquidity (COL). Thus, in order to compute daily $L - VaR$ estimates, we need to obtain daily VaR estimates through the models explained in the previous subsections, and daily COL estimates.

Recalling Equation 9, COL_t is based on P_t, the mid-price at time t; $\bar{S}$, the average relative bid-ask spread; $\tilde{\sigma}_t$, the volatility of the relative bid-ask spread at time t; and a, which is a scaling factor that produces $(1 - \alpha)\%$ coverage probability, where α is the significance level of the VaR being estimated. Since we will be estimating the VaR using $\alpha = 5\%$, a will be a scaling factor which produce 95% coverage probability. To obtain a series of daily COL_t estimates, the daily relative spread was computed, and its arithmetic mean was extrapolated as $\bar{S}$. Next, the relative spread volatility estimates, $\tilde{\sigma}$, were obtained using the same volatility models described in Section 3.1. Lastly, the scaling factor a was computed as the difference between the 95th percentile spread and the average spread of the empirical data, scaled by the average volatility of each model, which outputs a factor that adjusts for the spread's variation relative to volatility.

Both the scaling factor a and the mean relative spread $\bar{S}$ are thus constants in the daily estimation of the exogenous cost of liquidity at every time t. The value of a is dependent on the series of portfolio prices and specific volatility model being used, while $\bar{S}$ is solely reliant on the portfolio prices being used. Therefore, we expect the scaling factor a to be a scalar value that changes for each combination of portfolio returns and volatility models, whereas $\bar{S}$ will be a scalar value that changes solely depending on the portfolio being analysed. The only variables that shift daily and allow for a dynamic series of COL_t estimates are the volatility estimates $\tilde{\sigma}_t$ and the observed mid-prices P_t for each time t.

However, since we are computing the percentual VaR, the observed mid-price at time t, P_t, can be dropped from Equation 9, such that:

$$COL_t = \frac{1}{2}[(\bar{S} + a * \tilde{\sigma})] \tag{26}$$

where COL_t is now interpreted as a percentage of the total value of the portfolio being analysed.

3.3. Backtesting

After estimating each of the unadjusted VaR models described in Section 3.2. and each corresponding $L - VaR$ model, our objective is to evaluate the accuracy of each model. To do so, we will use the number of exceedances as our principal performance measure, where an exceedance is defined as an instance where the observed portfolio loss surpasses the VaR estimate for that same day (Alexander, 2009). For this purpose, we will conduct two well-established statistical tests: the Unconditional Coverage (UC) test proposed by Kupiec (1995) and the BCP test introduced by Berkowitz et al. (2011). Both tests provide different insights into model performance: the UC test will measure the observed exceedance rate of the estimates, while the BCP test will evaluate the presence of autocorrelation between exceedances.

From an exceedance rate perspective, we consider the models to perform properly when the actual exceedance rate is close to the expected exceedance rate for the given α significance level. The importance of maintaining an appropriate exceedance rate stems from a regulatory perspective, as most regulators rely on this metric to assess the capital requirements imposed on institutions (Jorion, 2006). A large exceedance rate can be proof that the market risk is being underestimated by the VaR models, which can cause regulators to increase the base capital requirements to hedge market risk. Hence, to remain within regulatory limits, the exceedance rate is an important tool to control to avoid excessive capital requirements.

Additionally, examining the autocorrelations of exceedances will help detect exceedance clustering, which can be an indicator that the data isn't adapting rapidly enough to harsh volatility moves in short periods of time. This isn't observable when solely looking at the exceedance rate, which simply measures the total amount of exceedances without considering the time at which they occur.

Even though both statistical tests provide different insights into the performance of the models, the exceedance rate remains the most useful indicator to measure since it holds significant regulatory importance and is a direct measure of risk, as it directly quantifies the frequency of model exceedances. As such, the main statistical test which will be used is the UC test, with the BCP test serving as an auxiliary tool to compare models which share similar UC test performances. This will be useful when choosing the most appropriate model, as even though a model can present an exceedance rate close to its expected exceedance rate, it can still fail the BCP test.

Both the UC test and BCP test are performed for all model combinations during the test period spanning from January 5, 2016, to December 29, 2023.

3.3.1. Unconditional Coverage test

Introduced by Kupiec (1995), the UC test is a likelihood ratio test that examines whether the number of exceedances observed in a sample matches the expected exceedances given by the confidence level α of the VaR model.

Formally, the indicator function $I_{\alpha,t}$ for each n observation in the sample can be defined as:

$$I_{\alpha,t+1} = \begin{cases} 1, & \text{if } r_{t+1} < -VaR_{1,\alpha,t} \\ 0, & \text{otherwise} \end{cases} \tag{27}$$

where r_t is the observed return at time t and $VaR_{1,\alpha,t-1}$ is the VaR estimate for day t.

The null hypothesis tests whether the indicator function, which is assumed to follow an i.i.d. Bernoulli process, presents a constant exceedance rate in line with the expected exceedance rate given the significance level α (Alexander, 2009). The null and alternative hypothesis are denoted as such:

$$\begin{gathered} H_0: \pi_{obs} = \pi_{exp} \equiv \alpha \\ H_a: \pi_{obs} \neq \pi_{exp} \end{gathered} \tag{28}$$

where π_{obs} is the observed exceedance rate in the sample and π_{exp} is the expected exceedance rate given the significance level α. The size of the sample is relevant to the accuracy of the UC test as the acceptance regions under the null hypothesis are larger for smaller sample sizes, which can lead to not rejecting the null even though it proves to be false. To overcome this, we will be using a sample size of 8 years to conduct the backtest.

The UC test statistic is given by:

$$LR_{UC} = \left(\frac{\pi_{exp}}{\pi_{obs}}\right)^{n_1} \left(\frac{1-\pi_{exp}}{1-\pi_{obs}}\right)^{n_0} \tag{29}$$

where n_1 is the number of exceedances and n_0 is the number of non-exceedances.

Under the null hypothesis, the asymptotic distribution of the test statistic is a chi-square with one degree of freedom: $-2\ln(LR_{UC}) \sim \chi^2_{(1)}$.

As we are estimating the daily VaR at the $\alpha\%$ significance level, we consider the models to be properly specified if we fail to reject the null hypothesis at the $(1-\alpha)\%$ confidence level, reflecting that the actual observed exceedance rate is in line with the expected α level exceedance rate.

3.3.2. BCP test

The BCP test (Berkowitz et al., 2011) is a Ljung-Box test that assesses the independence of the exceedances by checking for their first K autocorrelations. If the VaR model is well specified, we expect the autocorrelations of the exceedances to be 0 for all lags, demonstrating that we cannot predict when the next exceedance will occur. Hence, the null and alternative hypothesis are defined as:

$$\begin{gathered} H_0: \hat{\rho}_k = 0, \forall \mathrm{k} \in \{1, \dots, K\} \\ H_a: \exists k \in \{1, \dots, K\} \text{ s.t. } \hat{\rho}_k \neq 0 \end{gathered} \tag{30}$$

where $\hat{\rho}_k$ is the lag k sample autocorrelation of the output of the indicator function $I_{\alpha,t}$ described in Equation 27, and K is the maximum autocorrelation lag in the test.

The test statistic of the BCP test is given by:

$$BCP(K) = T(T+2)\sum_{k=1}^{K} \frac{\hat{\rho}_k^2}{T-k} \tag{31}$$

where T is the sample size of the test.

Under the null hypothesis, the asymptotic distribution of the BCP test statistic is a chi-square with K degrees of freedom: $BCP(K) \sim \chi^2_{(K)}$.

The choice of the lag K for backtesting is dependent on the user. There exists a compromise between the choice of a smaller and larger K: a larger K aids in detecting non-independence at higher lags, but at the cost of a decrease in the power of the test. As the BCP test statistic follows a chi-square with K degrees of freedom, increasing the number of degrees of freedom causes the distribution to become more spread out, decreasing the non-acceptance area and making the null hypothesis harder to reject. Therefore, we will begin with $K = 1$ and increase it until $K = 5$, to test for autocorrelation until the 5th lag.

4. Model selection and cost of liquidity analysis

In this section we present the results of our empirical analysis. We present the validation of the estimates produced by each model for each portfolio, using the UC and BCP tests to evaluate their performance for the test period. Our goal will be to assess whether incorporating the cost of liquidity into the VaR improves the performance of each model.

The test period is the 8-year period from January 5, 2016, to December 29, 2023, which constitutes a total number of observations of $n = 2011$. Taking into account our significance level of $\alpha = 5\%$, we will fail to reject the null hypothesis of the UC test if the total number of exceedances falls close to 100 exceedances. If the *p*-value of a model falls below 5%, we will reject the null hypothesis of the UC test, and the model is deemed to be poorly specified. For the models that present a *p*-value higher than 5%, and are thus accepted by the UC test, we will rely on the BCP test to compare which model is the best specified.

We will first analyze the standard VaR models without considering the liquidity component, and then perform the same tests after applying the cost of liquidity to the VaR estimates, to understand whether incorporating the liquidity risk adds validity to the models.

Table 4.1 presents each model that will be analyzed per portfolio, and its respective assigned number for ease of comparison.

Table 4.1

Model description and respective numbering

Model Nº	Description
1	T+1 Volatility-adjusted Historical, 252 rolling observations, GARCH volatility
2	T+1 Volatility-adjusted Historical, 252 rolling observations, student t GARCH volatility
3	T+1 Volatility-adjusted Historical, 252 rolling observations, FIGARCH volatility
4	Monte Carlo simulation, 252 rolling observations, GARCH volatility
5	Monte Carlo simulation, 252 rolling observations, student t GARCH volatility
6	Monte Carlo simulation, 252 rolling observations, FIGARCH volatility
7	Parametric normal, GARCH volatility
8	Parametric normal, student t GARCH volatility
9	Parametric normal, FIGARCH volatility

Note. Includes a brief description of each of the VaR models estimated, with a corresponding labeled number for ease of distinction and comparison.

4.1. Backtesting illiquid portfolio

In Table 4.2 below we present the summary of the backtest for the non-adjusted VaR and the liquidity adjusted VaR of the illiquid portfolio, which includes the number of exceedances per model, the exceedance rate, and the *p*-value of the UC test. If the *p*-value falls below 5%, the UC test null hypothesis is rejected, and the model is considered poorly specified. As such, the values in bold denote the models which present a *p*-value above or equal to 5%, i.e., models that are accepted by the UC test.

Table 4.2

Summary of the backtest of the illiquid portfolio

Model Class	Model Nº	Non-adjusted VaR			Liquidity-adjusted VaR		
		Nº Exceedances	Exc. Rate (%)	*p*-value (%)	Nº Exceedances	Exc. Rate (%)	*p*-value (%)
Historical	1	102	5.07	**88.23**	97	4.82	**71.49**
	2	102	5.07	**88.23**	95	4.72	**56.67**
	3	110	5.47	**34.05**	99	4.92	**87.37**
Monte Carlo	4	119	5.92	**6.62**	111	5.52	**29.26**
	5	125	6.22	1.58	115	5.72	**14.79**
	6	123	6.12	2.62	110	5.47	**34.05**
Parametric Normal	7	106	5.27	**58.03**	97	4.82	**71.49**
	8	105	5.22	**65.11**	91	4.53	**32.10**
	9	108	5.37	**45.11**	100	4.97	**95.51**

Note. For the illiquid portfolio we present the exceedance rate and the *p*-value of the UC test of each model. Highlighted in bold are the models accepted by the UC test, with *p*-values higher than 5%. Refer to Table 4.1 for the description of each model.

Looking at the results of the non-adjusted VaR, almost all models pass the UC test, with the Monte Carlo models presenting the poorest results. Out of all 3 Monte Carlo models, only model 4, the model with standard GARCH volatility, passes the test with a *p*-value of 6.62%.

Observing the results of the liquidity adjusted VaR, every model passes the UC test. Models 5 and 6, which were previously rejected, are now accepted after incorporating the liquidity component, proving that when the cost of liquidity is incorporated, the total number of exceedances can decrease enough to add validity to a model.

Overall, the models with the best UC test results for the non-adjusted VaR of the illiquid portfolio are the volatility-adjusted historical VaR models. For the liquidity adjusted models, model 9 presents the highest UC test *p*-value of 95.51%. After incorporating the liquidity component into the VaR estimation of model 9, its UC test p-value increased from 45.11% to 95.51%, substantially improving its accuracy.

It may be interesting to note how the number of exceedances decreases when the cost of liquidity is considered. Since the illiquid portfolio is composed of less liquid assets, we expect the cost of liquidity to have a higher impact on the liquidity adjusted VaR when compared to the liquid portfolio. Table 4.3 shows the number of exceedances per model, for both the non-adjusted and the liquidity adjusted VaR, as well as the difference of exceedances when the cost of liquidity is included in the VaR calculation of the illiquid portfolio.

Table 4.3

Comparison of exceedances between non-adjusted VaR and $L - VaR$ of the illiquid portfolio

		Non-adjusted VaR	Liquidity-adjusted VaR	
Model Class	Model Nº	Nº Exceedances	Nº Exceedances	Exc. Difference
Historical	1	102	97	5
	2	102	95	7
	3	110	99	11
Monte Carlo	4	119	111	8
	5	125	115	10
	6	123	110	13
Parametric Normal	7	106	97	9
	8	105	91	14
	9	108	100	8

Note. We display the difference in exceedances that arises from adding the cost of liquidity to the non-adjusted VaR models for the illiquid portfolio. Refer to Table 4.1 for the description of each model.

As shown in Table 4.3, considering the cost of liquidity for the illiquid portfolio causes a significant decrease in the number of exceedances of the models, proving that the liquidity cost is a significant risk that should be modelled and accounted for when hedging the total portfolio risk. Models 6 and 8, for example, see a decrease of 13 and 14 exceedances, respectively.

Since every single $L - VaR$ model for the illiquid portfolio passes the UC test, we will use the BCP test to determine which model performs the best. Table 4.4 summarizes the results of the BCP test.

Table 4.4

Illiquid portfolio - BCP test results

Model Class	Model Nº	Standard VaR					Liquidity adjusted VaR				
		Lag 1	Lag 2	Lag 3	Lag 4	Lag 5	Lag 1	Lag 2	Lag 3	Lag 4	Lag 5
Historical	1	0.03	0.00	0.00	0.00	0.00	0.04	0.00	0.00	0.00	0.00
	2	0.00	0.00	0.00	0.00	0.00	0.02	0.00	0.00	0.00	0.00
	3	0.01	0.00	0.00	0.00	0.00	0.07	0.00	0.00	0.00	0.00
Monte Carlo	4	**11.29**	3.97	**8.51**	**12.23**	**12.03**	**21.94**	2.01	4.66	**7.04**	**6.92**
	5	**10.57**	0.59	0.47	0.35	0.78	2.49	0.24	0.27	0.27	0.60
	6	**8.21**	2.30	**5.54**	**9.38**	**15.77**	3.16	0.36	0.74	1.61	3.24
Parametric Normal	7	4.87	**11.13**	**18.14**	**19.26**	**25.94**	0.97	3.34	**7.54**	**8.37**	**12.32**
	8	1.29	2.42	**5.80**	3.90	**7.17**	0.24	0.84	2.27	3.20	**6.08**
	9	2.25	**7.37**	**14.94**	**17.74**	**13.79**	**5.73**	**16.42**	**30.62**	**42.37**	**43.65**

Note. Showcases the p-values (%) of the BCP test at all lags for each model of the illiquid portfolio that passed the UC test. Highlighted in bold are the models accepted by the BCP test, with a p-value higher than 5%. Refer to Table 4.1 for the description of each model.

For the non-adjusted VaR, no single model passes the BCP test at all lags. However, models 5, 6, and 7 all pass the BCP test at 4 lags. For the liquidity adjusted VaR, the model with the best performance under the BCP test is model 9, which now manages to pass the test at all lags. Comparing the non-adjusted and liquidity adjusted lag 1 p-value of model 9, it seems that including the cost of liquidity in the VaR estimation lightens the existence of consecutive exceedances, since the p-value increases and model 9 passes the BCP test for lag 1. Similarly, the p-value of lag 1 for model 4 also increases when the cost of liquidity is included.

With regards to models 5, 6, 7 and 8, including the cost of liquidity in the VaR estimation highlights an intriguing dynamic of the exceedance autocorrelations. For these models, when the liquidity component is considered, the number of exceedances decreases. However, the decrease in exceedances results in an increase of the autocorrelations at certain lags. Observing model 6, when the cost of liquidity is added to the VaR, the p-value of the autocorrelations at all lags significantly decreases – most notably at lag 5, where the non-adjusted VaR model passes the BCP test, but the $L - VaR$ at the same lag fails the test.

This is likely being caused by uncaptured dynamics of the models' risks. When the exceedances decrease, we observe a slight increase in their clustering, indicating that the exceptions are becoming more predictable over time.

Comparing the BCP test results of the non-adjusted VaR models, models 5, 6 and 7 present the best test results. However, models 5 and 6 are rejected by the UC test, and thus we choose model 7 as the best non-adjusted VaR model for the illiquid portfolio, as it has a UC test *p*-value of 58.03%. Considering that the exceedance rate will be our main metric of measurement for a model's validity, it is worthy to note that the volatility-adjusted historical VaR models present the highest *p*-values of the UC test amongst the non-adjusted models. Nonetheless, they all fail the BCP test, presenting *p*-values no higher than 1%. As such, model 7 is considered the best non-adjusted VaR model for the illiquid portfolio.

For the liquidity adjusted VaR, model 9 is the model with the best test performance, as it presents a UC test *p*-value of 95.51% and passes the BCP test at all lags.

4.2. Backtesting liquid portfolio

Following the same analysis done for the illiquid portfolio, in Table 4.5 we present the backtest results of the non-adjusted VaR and $L - VaR$ of the liquid portfolio. The table includes the number of exceedances per model, the exceedance rate, and the *p*-value of the UC test. As per our previous analysis of the illiquid portfolio, the models accepted by the UC test are denoted by the values in bold.

Table 4.5

Summary of the backtest of the liquid portfolio

Model Class	Model Nº	Non-adjusted VaR			Liquidity adjusted VaR		
		Nº Exceedances	Exc. Rate (%)	*p*-value (%)	Nº Exceedances	Exc. Rate (%)	*p*-value (%)
Historical	1	115	5.72	**14.79**	110	5.47	**34.05**
	2	113	5.62	**21.12**	110	5.47	**34.05**
	3	120	5.97	**5.31**	116	5.77	**12.24**
Monte Carlo	4	103	5.12	**80.28**	98	4.87	**79.33**
	5	105	5.22	**65.11**	98	4.87	**79.33**
	6	93	4.62	**43.42**	92	4.57	**37.51**
Parametric Normal	7	87	4.33	**15.63**	81	4.03	3.87
	8	92	4.57	**37.51**	88	4.38	**18.99**
	9	80	3.98	2.95	76	3.78	0.88

Note. For the liquid portfolio we present the exceedance rate and the *p*-value of the UC test of each model. Highlighted in bold are the models accepted by the UC test, with *p*-values higher than 5%. Refer to Table 4.1 for the description of each model.

With regards to the liquid portfolio, the only non-adjusted model which is rejected by the UC test is model 9, the parametric normal VaR model with FIGARCH volatility. Out of all the other models, the one with the highest UC test *p*-value is model 4, the Monte Carlo model with GARCH volatility. Observing the liquidity adjusted models, model 4 remains the model with the highest UC test *p*-value, now alongside model 5, which shares the same *p*-value.

It is important to note that model 9, the model that is being rejected by the UC test, presents a lower exceedance rate than the models accepted by the UC test. The UC test functionally rejects not only models which are underestimated, and thus have a larger number of exceedances than expected, but also rejects models which are overestimated, with very few exceedances. In terms of capital allocation requirements, models with a large number of exceedances cause an underestimation of the relevant risk factors, which causes insufficient capital reserves being held by institutions, and may lead to non-compliance of the regulatory frameworks. In contrast, models with few exceedances cause an overestimation of the risk factors, and the excessive capital requirements that stem from these more conservative estimates can reduce operational efficiency, as well as cause difficulty in distinguishing which risks are being properly hedged. In our analysis of the liquid portfolio, the models which are being rejected by the UC test are so because of an exceedance overestimation. As such, models which were already rejected by the UC test due to being overestimated will still be rejected when the cost of liquidity is incorporated, as the number of exceedances is expected to become even more conservative if we include the added exogenous cost of liquidity.

Indeed, observing Table 4.5, the liquidity adjusted model 9 is also rejected by the UC test due to an exceedance overestimation. Furthermore, model 7 is also rejected when the cost of liquidity is incorporated due to the same reason.

Table 4.6 presents the difference of exceedances between the non-adjusted VaR and the $L-VaR$ models of the liquid portfolio, to understand how adding the cost of liquidity to the VaR estimate reduces the number of exceedances. As shown in Table 4.6, the cost of liquidity for the liquid portfolio contributes significantly less to overall risk compared to the illiquid portfolio. This result is expected since liquid assets have a larger trade volume than illiquid assets. The large number of market players willing to buy or sell these assets distributes the liquidity risk across the market, thereby reducing its impact.

Table 4.6

Comparison of exceedances between standard VaR and $L - VaR$ of the liquid portfolio

		Non-adjusted VaR	Liquidity adjusted VaR	
Model Class	Model Nº	Nº Exceedances	Nº Exceedances	Exc. Difference
Historical	1	115	110	5
	2	113	110	3
	3	120	116	4
Monte Carlo	4	103	98	5
	5	105	98	7
	6	93	92	1
Parametric Normal	7	87	81	6
	8	92	88	4
	9	80	76	4

Note. We display the difference in exceedances that arises from adding the cost of liquidity to the non-adjusted VaR models for the liquid portfolio. Refer to Table 4.1 for the description of each model.

Table 4.7 shows a summary of the BCP test results of the liquid portfolio. The table presents the *p*-value for all lags of the models which passed the UC test.

Table 4.7

Liquid portfolio - BCP test results

Model Class	Model Nº	Standard VaR					Liquidity adjusted VaR				
		Lag 1	Lag 2	Lag 3	Lag 4	Lag 5	Lag 1	Lag 2	Lag 3	Lag 4	Lag 5
Historical	1	**6.73**	0.00	0.00	0.00	0.00	3.16	0.00	0.00	0.00	0.00
	2	**5.05**	0.00	0.00	0.00	0.00	3.16	0.00	0.00	0.00	0.00
	3	**12.71**	0.00	0.00	0.00	0.00	**7.71**	0.00	0.00	0.00	0.00
Monte Carlo	4	**21.13**	0.39	1.06	1.83	2.81	**28.48**	2.40	**5.81**	**6.96**	**10.70**
	5	4.17	0.17	0.50	0.97	1.99	**12.09**	1.28	3.30	4.14	**6.68**
	6	**72.43**	**5.58**	**8.76**	**7.51**	**12.46**	**68.67**	4.61	**9.63**	**7.63**	**12.46**
Parametric Normal	8	0.31	1.25	2.99	2.58	4.99	0.60	2.30	**5.62**	3.32	**6.26**

Note. Showcases the *p*-values (%) of the BCP test at all lags for each model of the liquid portfolio that passed the UC test. Highlighted in bold are the models accepted by the BCP test, with a *p*-value higher than 5%. Refer to Table 4.1 for the description of each model.

Among the non-adjusted models, the model with the best BCP test results is model 6, which passes the test at all lags. For the liquidity adjusted models, both models 4 and 6 pass the BCP test for all lags except lag 2, with model 6 presenting bigger *p*-values for all passing lags.

Overall, for the non-adjusted models, model 6 presents the best backtest performance, as it is accepted by the BCP test at all lags and isn't rejected by the UC test. For the liquidity adjusted models, both models 4 and 6 pass the BCP test for all lags expect lag 2, with model 6 slightly outperforming model 4 at every passing lag. However, model 4 presents a UC test *p*-value of 79.33%, which is higher than model 6's *p*-value of 37.51%. As mentioned in Section 3.3., the exceedance rate is the most relevant metric through which we can evaluate a model's validity, while the BCP test would be used as an auxiliary tool to compare models with similar UC test results. Since model 4 presents a much higher UC test *p*-value, we consider model 4 to be the best liquidity-adjusted VaR model, and model 6 as the best non-adjusted VaR model due to its better performance in the BCP test.

4.3. Backtesting conclusions

Table 4.8 presents a summary for each portfolio of the models with the best backtest performance.

Table 4.8

Backtest result summary for all models analyzed

	Illiquid portfolio		Liquid portfolio	
	Non-adjusted VaR	$L-VaR$	Non-adjusted VaR	$L-VaR$
VaR Models	Parametric Normal - GARCH	Parametric Normal - FIGARCH	Monte Carlo - FIGARCH	Monte Carlo - GARCH

Note. Highlights the models with the best backest results for each portfolio, distinguishing between non-adjusted VaR and L-VaR.

Regarding the liquid portfolio, the models that most properly capture the risk dynamics of the underlying returns are models 6 for the non-adjusted VaR, and model 4 for the liquidity adjusted VaR.

With respect to the illiquid portfolio, model 7 is considered as the best model amongst the non-adjusted VaR models. However, it is outperformed by model 9 when the liquidity component is considered.

As shown in Tables 4.3 and 4.6, there is a stark distinction between the impact the cost of liquidity holds when dealing with illiquid assets. When the liquidity cost was incorporated into the VaR measures of the illiquid portfolio, the number of exceedances greatly decreased, showing that the standard VaR is uncapable of capturing the true underlying dynamics of less liquid assets. On the other hand, when the liquidity cost was incorporated into the VaR estimates of the liquid portfolio, there wasn't a notable decrease in the number of exceedances, indicating that considering this cost didn't provide substantial improvements to the non-adjusted VaR. On average, incorporating the cost of liquidity for the illiquid portfolio causes a decrease of 9 exceedances per model, while incorporating the cost of liquidity for the liquid portfolio causes a decrease of 4 exceedances per model.

Even though the $L - VaR$ measures can start to deviate from the expected exceedance rate as they become more conservative, including the exogenous cost of liquidity is shown to substantially decrease the number of exceedances for less liquid assets. This demonstrates that liquidity risk is a major component of illiquid assets' total risk, which should be properly modelled to fully capture their risk dynamics.

Conclusion

In this study, we explored whether Value-at-Risk estimation could be improved by incorporating the added cost of exogenous liquidity risk. Since the Value-at-Risk is a statistical estimate that only measures market risk, we expect liquidity risk to be prominent for less liquid assets, and as such modelling it to be necessary in order to fully capture the true risk dynamics that an investor or institution faces when holding these assets.

Two portfolios were studied in this work, one comprised of less liquid U.S. equity stocks, and another comprised of more liquid U.S. equity stocks, chosen based on their intraday traded volume. The estimates for the VaR were computed using three different volatility models, GARCH, student t GARCH and FIGARCH. To diversify our results, we estimated the VaR using three different methods, the parametric normal VaR, the volatility-adjusted historical VaR, and the Monte Carlo VaR. Each model was backtested using the UC and BCP tests.

The backtest performed on all models presented expected results. When comparing the non-adjusted VaR and the liquidity adjusted VaR, there was a substantial difference in the reduction of exceedances observed between the liquid and the illiquid portfolio. Since liquid assets are less exposed to liquidity issues, we expect the non-adjusted Value-at-Risk to correctly model most of the risks faced by the investors or institutions that hold these assets. Contrarily, we expect the VaR of the illiquid portfolio to be underestimated by non-adjusted VaR models due to their illiquidity, which comprises a large part of the risk one is exposed to whilst holding these assets.

For the illiquid portfolio, non-adjusted VaR models were shown to underestimate the actual risk, presenting an observed exceedance rate higher than the expected rate. When including the cost of liquidity, some models that failed to pass the UC test due to their large number of exceedances now presented an exceedance rate much more in accordance with the expected exceedance rate of 5%. As such, we conclude that adding the cost of liquidity to the VaR more accurately captures the true risk dynamics of illiquid assets.

For the liquid portfolio, incorporating the liquidity component didn’t provide an increase in the accuracy of the models. Most models were accepted by the UC test and including the liquidity component did not decrease the number of exceedances enough to warrant an increase in accuracy.

Overall, the parametric normal VaR models were more successful in modelling the risk of the illiquid portfolio, while the Monte Carlo models were more successful in modelling the risk of the liquid portfolio. On average, incorporating the cost of liquidity for the illiquid portfolio caused a decrease of 9 exceedances per model, compared to the liquid portfolio's average of 4, which showcases the importance of measuring liquidity risk for less liquid assets.

For future research, it would be interesting to delve deeper into the behavior of the exceedances' autocorrelations. For some liquidity adjusted models, the BCP test was rejected at the same lags where their non-adjusted counterpart was not rejected. Understanding this behavior would aid in highlighting the full impact that adding the cost of liquidity has on the VaR. Furthermore, it would also be interesting to incorporate the cost of liquidity into other types of assets such as bonds, which are largely dependent on their liquidity for their pay-off.

Bibliography

Affleck-Greaves, J., Callahan, C., & Ramachandran, R. (2000). Detecting Abnormal Bid-Ask Spread: A Comparison of Event Study Methods. *Review of Quantitative Finance and Accounting, 14(1),* 45-65.

Al Janabi, M. A. (2008). Integrating Liquidity Risk Factor into a Parametric Value at Risk Method. *The Journal of Trading, 3*(3), 76-87.

Alexander, C. (2009). *Market Risk Analysis, Value at Risk Models Volume IV.* John Wiley & Sons.

Allen, F., Goldstein, I., & Jagtiani, J. (2018). The Interplay among Financial Regulations, Resilience, and Growth. *Journal of Financial Services Research, 53,* 141-162.

Baillie, R. T., Bollerslev, T., & Mikkelsen, H. O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics, 74*(1), 3-30.

Bangia, A., Diebold, F. X., Schuermann, T., & Stroughair, J. D. (2008, May). Modeling Liquidity Risk with Implications for Traditional Market Risk Measurement and Management. *NYU Working Paper No. FIN-99-062*.

Berkowitz, J. (2000). Incorporating Liquidity Risk into Value-at-Risk Models. *The Journal of Derivatives*.

Berkowitz, J., Christoffersen, P., & Pelletier, D. (2011). Evaluating Value-at-Risk Models with Desk-Level Data. *Management Science*.

Bertsimas, D., & Lo, A. W. (1998). Optimal Control of Liquidation Costs. *Journal of Financial Markets, 1*(1), 1-50.

Black, F. (1971). Towards a Fully Automated Stock Exchange. *Financial Analysts Journal, Vol. 27,* 28-35.

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

Bollerslev, T. (1987). A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return. *The Review of Economics and Statistics, 69*(3), 542-547.

Boothe, P., & Glassman, D. (1987). The Statistical Distribution of Exchange Rates: Empirical Evidence and Economic Implications. *Journal of International Economics, 22*(3-4), 297-319.

Bowman, K., & Shenton, L. (1975). Omnibus Contours for Departures from Normality Based on √b1 and b2. *Biometrika, 62,* 243-250.

Bradley, B. O., & Taqqu, M. S. (2003). Financial Risk and Heavy Tails. In S. Rachev, *Handbook of Heavy Tailed Distributions in Finance* (pp. 35-103). North Holland.

Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica, 50*(4), 987-1007.

Engle, R. F., & Bollerslev, T. (1986). Modelling the persistence of conditional variances. *Econometric Reviews, 5*(1), 1-50.

Fama, E. F. (1965). The Behavior of Stock-Market Prices. *Journal of Business*, 34-105.

Fielitz, B. D., & Rozelle, J. P. (1983). Stable Distributions and the Mixture of Distributions Hypothesis for Common Stock Returns. *Journal of the American Statistical Association, 78*, 28-36.

Hansen, P. R., & Lunde, A. (2005). A Forecast Comparison of Volatility Models: Does Anything Beat a Garch (1,1)? *Journal of Applied Econometrics, 20*(7), 873-889.

Hull, J., & White, A. (1998). Incorporating volatility updating into the historical simulation method for VaR. *The Journal of Risk*, *1*(1), 5-19.

J.P. Morgan and Reuters. (1996). Riskmetrics - technical document. Retrieved from https://www.msci.com/documents/10199/5915b101-4206-4ba0-aee2-3449d5c7e95a

Jarque, C. M., & Bera, A. K. (1987). A Test for Normality of Observations and Regression Residuals. *Revue Internationale de Statistique, Vol. 55, No. 2*, 163-172.

Johnson, N., Kotz, S., & Balakrishnan, N. (1995). *Continuous Univariate Distributions, Vol. 2, Chapter 28.* New York: J. Wiley & Sons.

Jorion, P. (2006). *Value at Risk: The New Benchmark for Managing Financial Risk, 3rd Edition.* McGraw Hill.

Kim, J., Malz, A., & Mina, J. (1999). LongRun Technical Document, RiskMetrics Group. New York.

Kupiec, P. (1995). Techniques for Verifying the Accuracy of Risk Measurement Models. *Journal of Derivatives, 3*, 73-84.

Kyle, A. S. (1985). Trading, Continuous Auctions and Insider. *Econometrica, Vol. 53*, 1315-1335.

Shapiro, S., & Wilk, M. (1965). An Analysis of Variance Test for Normality (Complete Samples). *Biometrika, Vol. 52, No. 3/4*, 591-611.

Annexes

Annex A – Source Code

```
### Required Libraries ###

import pandas as pd

import numpy as np

from scipy.stats import skew, kurtosis, jarque_bera, shapiro, norm, chi2

from arch import arch_model


### Function to obtain descriptive statistics and normality test results of the portfolio returns and spreads ###

def analyze_portfolio(returns, name):
    return_stats = returns.describe()
    if isinstance(return_stats, pd.DataFrame):
        mean_return = return_stats.loc['mean', 'Returns']
        median_return = return_stats.loc['50%', 'Returns']
        max_return = return_stats.loc['max', 'Returns']
        std_return = return_stats.loc['std', 'Returns']
    elif isinstance(return_stats, pd.Series):
        mean_return = return_stats.loc['mean']
        median_return = return_stats.loc['50%']
        max_return = return_stats.loc['max']
        std_return = return_stats.loc['std']
    else:
        raise ValueError("Unexpected return_stats type")
    statistics = pd.DataFrame({
        'Mean': [mean_return],
        'Median': [median_return],
        'Max': [max_return],
        'Standard Deviation': [std_return],
        'Skewness': [skew(returns)],
        'Kurtosis': [kurtosis(returns)]
    }, index=[f'{name} Returns'])
```

```
    ### Jarque-Bera test for normality ###
    jb_stat, jb_pvalue = jarque_bera(returns)
    jb_test_result = 'Normal' if jb_pvalue > 0.05 else 'Non-Normal'
    ### Shapiro-Wilk test for normality ###
    sw_stat, sw_pvalue = shapiro(returns)
    sw_test_result = 'Normal' if sw_pvalue > 0.05 else 'Non-Normal'
    normality_test_results = pd.DataFrame({
        'Test': ['Jarque-Bera', 'Shapiro-Wilk'],
        'Test Statistic': [jb_stat, sw_stat],
        'P-value': [jb_pvalue, sw_pvalue],
        'Result': [jb_test_result, sw_test_result]
    })
    return statistics, normality_test_results

### Function to obtain the volatility estimates ###
def compute_volatilities(returns, window_size=252, forecast_horizon=21):
    num_forecasts = len(returns) - window_size
    garch_volatilities = []
    t_garch_volatilities = []
    figarch_volatilities = []
    for i in range(num_forecasts):
        window_data = returns[i:i + window_size]
        ### GARCH(1,1) with Normal distribution ###
        model_garch = arch_model(window_data, vol="GARCH", p=1, q=1)
        results_garch = model_garch.fit(disp='off', options={'maxiter': 5000, 'ftol': 1e-5})
        forecast_garch = results_garch.forecast(horizon=forecast_horizon)
        variance_garch = forecast_garch.variance.iloc[-1, 0]
        volatility_garch = np.sqrt(variance_garch)
        garch_volatilities.append(volatility_garch)
        ### GARCH(1,1) with Student's t distribution ###
        model_t_garch = arch_model(window_data, vol="GARCH", p=1, q=1, dist='t')
```

```
        results_t_garch = model_t_garch.fit(disp='off', options={'maxiter': 5000, 'ftol': 1e-5})
        forecast_t_garch = results_t_garch.forecast(horizon=forecast_horizon)
        variance_t_garch = forecast_t_garch.variance.iloc[-1, 0]
        volatility_t_garch = np.sqrt(variance_t_garch)
        t_garch_volatilities.append(volatility_t_garch)
        ### FIGARCH(1,d,1) ###
        model_figarch = arch_model(window_data, vol="FIGARCH", p=1, q=1)
        results_figarch = model_figarch.fit(disp='off', options={'maxiter': 5000, 'ftol': 1e-5})
        forecast_figarch = results_figarch.forecast(horizon=forecast_horizon)
        variance_figarch = forecast_figarch.variance.iloc[-1, 0]
        volatility_figarch = np.sqrt(variance_figarch)
        figarch_volatilities.append(volatility_figarch)
    volatility_df = pd.DataFrame({
        'GARCH_volatility': garch_volatilities,
        't_GARCH_volatility': t_garch_volatilities,
        'FIGARCH_volatility': figarch_volatilities
    }, index=returns.index[window_size:])
    return volatility_df

### Function to compute volatility-adjusted returns ###
def vol_adjusted_returns(returns, volatilities):
    adjusted_returns = pd.DataFrame(index=returns.index)
    for model in volatilities.columns:
        model_adjusted_returns = pd.Series(index=returns.index)
        T = len(returns)
        for t in range(T-1):
            sigma_t = volatilities.iloc[t][model]
            sigma_T = volatilities.iloc[t+1][model]
            model_adjusted_returns.iloc[t] = (sigma_T / sigma_t) * returns.iloc[t]
        adjusted_returns[model] = model_adjusted_returns
    return adjusted_returns.dropna()
```

```
### Function to obtain historical VaR estimates ###
def historical_var(returns, training_size=252, alpha=0.05):
    historical_var_results = []
    for i in range(training_size, len(returns)):
        returns_subset = returns.values[:i]
        var_result = -np.percentile(returns_subset, 100 * alpha, axis=0)
        historical_var_results.append(var_result)
        var_df = pd.DataFrame(historical_var_results, columns=returns.columns, index=returns.index[training_size:])
    return var_df.dropna()


### Function to obtain parametric normal VaR estimates ###
def parametric_var(volatilities, confidence_level=0.95):
    score = norm.ppf(confidence_level)
    parametric_var = volatilities * score
    return parametric_var


### Function to obtain Monte Carlo VaR estimates ###
def rolling_monte_carlo_var_with_volatility(portfolio_returns, volatilities, window_size=252, num_simulations=1000, confidence_level=0.95):
    if not isinstance(portfolio_returns, pd.DataFrame):
        raise ValueError("portfolio_returns should be a pandas DataFrame.")
    if not isinstance(volatilities, pd.DataFrame):
        raise ValueError("volatilities should be a pandas DataFrame.")
    var_series = pd.DataFrame(index=portfolio_returns.index[window_size:], columns=portfolio_returns.columns, dtype=float)
    for model in portfolio_returns.columns:
        for i in range(window_size, len(portfolio_returns)):
            window_returns = portfolio_returns[model].iloc[i - window_size:i]
            window_volatilities = volatilities[model].iloc[i - window_size:i]
            mean_return = np.mean(window_returns)
            model_volatilities = window_volatilities.values
```

```
            simulated_returns = np.zeros((num_simulations, window_size))
            for j in range(num_simulations):
                for k in range(window_size):
                    sigma = model_volatilities[k]  # Volatility for the current day
                    simulated_returns[j, k] = np.random.normal(mean_return, sigma)
            initial_portfolio_value = 1
            simulated_portfolio_values = initial_portfolio_value * (1 + simulated_returns).cumprod(axis=1)
            simulated_portfolio_returns = np.diff(simulated_portfolio_values, axis=1) / simulated_portfolio_values[:, :-1]
            var_percentile = np.percentile(simulated_portfolio_returns, (1 - confidence_level) * 100, axis=0)
            VaR = -var_percentile[-1]  # VaR is typically reported as a positive number
            var_series.loc[portfolio_returns.index[i], model] = VaR
    return var_series


### Function to compute the cost of liquidity ###
def COL(data, vol, alpha=95):
    def spread_percentile(data):
        return np.percentile(data, alpha)
    percentile_spread = spread_percentile(data)
    av_data_spread = data.mean()
    a = {}
    for model in vol.columns:
        try:
            model_vol = vol[model].iloc[:-1].mean()
            a[model] = (percentile_spread - av_data_spread) / model_vol
            print(f"Computed a for {model}: {a[model]}")
        except ZeroDivisionError:
            print(f"ZeroDivisionError for model {model}")
            a[model] = np.nan
    col = pd.DataFrame(index=vol.iloc[:-1].index)
    for model in vol.columns:
```

```
        col[model] = 0.5 * (av_data_spread + a[model] * vol[model].iloc[:-1])
    return col

### Function to obtain liquidity adjusted VaR ###
def liquidity_adjusted_var(var, col):
    var = var.reindex(col.index)
    shifted_col = col.shift(1)
    shifted_col.iloc[-1] = col.iloc[-2]
    Lvar = var + shifted_col
    Lvar = Lvar.dropna()
    return Lvar

### Function for the UC and BCP test ###
### UC test ###
def unconditional_coverage_test(violations, p=0.05):
    T = len(violations)
    N = np.sum(violations)
    pi_hat = N / T
    L0 = (1 - p) ** (T - N) * p ** N
    L1 = (1 - pi_hat) ** (T - N) * pi_hat ** N
    LR_uc = -2 * np.log(L0 / L1)
    critical_value_uc = chi2.ppf(0.95, 1)
    decision = "Reject" if LR_uc > critical_value_uc else "Do not reject"
    return {
        'Test Statistic': LR_uc,
        'Critical Value': critical_value_uc,
        'Decision': decision
    }
### BCP test ###
def independence_test_lags(violations, max_lag=5):
    T = len(violations)
    V = np.array(violations)
```

```
results = {}
for lag in range(1, max_lag + 1):
    V_lag = np.roll(V, lag)
    V_lag[:lag] = 0
    N00 = np.sum((V == 0) & (V_lag == 0))
    N01 = np.sum((V == 1) & (V_lag == 0))
    N10 = np.sum((V == 0) & (V_lag == 1))
    N11 = np.sum((V == 1) & (V_lag == 1))
    p01 = N01 / (N00 + N01)
    p11 = N11 / (N10 + N11)
    L0 = ((1 - p01) ** N00) * (p01 ** N01) * ((1 - p01) ** N10) * (p01 ** N11)
    L1 = ((1 - p01) ** N00) * (p01 ** N01) * ((1 - p11) ** N10) * (p11 ** N11)
    LR_ind = -2 * np.log(L0 / L1)
    critical_value_ind = chi2.ppf(0.99, 1)
    decision = "Reject" if LR_ind > critical_value_ind else "Do not reject"
    results[f'Lag {lag}'] = {
        'Test Statistic': LR_ind,
        'Critical Value': critical_value_ind,
        'Decision': decision
    }
return results
```