Università LUISS Guido Carli

Department of Business and Management

Master of Science in Corporate Finance

Thesis in Risk Management

Value at Risk: an Experimental Analysis of Different Approaches

The supervisor

Prof. Daniele Penza

The co-supervisor

Prof. Nicola Borri

The Student

Michele Addamiano

Academic year 2021/2022

Contents

1 Introduction

The main objective of this work is to present and discuss the results of the analysis that I carried out to compare different approaches to Value at Risk (VaR), one of the main tools that banks adopt for risk management purposes.

With the objective of preserving financial stability, regulators in most areas of the world have decided to impose that banks at all times have layers of capital that work as a buffer in case banks face unexpected losses, and VaR is the measure that most financial institution have been adopting in the last decades to estimate the level of capital they must put aside.

The structure of the thesis is organized as follows: three chapters that follow this brief introduction and that constitute the main body of the dissertation, and a final section containing the conclusions.

The first chapter contains a theoretical analysis of what is Value at Risk. Here I begin with a general overview of banks risk management practices with a focus on market risk. Then, I proceed with a deep dive on the VaR concept, its merits and its drawbacks and the regulatory foundations on which VaR is rooted. I then present the three standard approaches adopted to estimate VaR, namely the Variance-Covariance approach, the historical simulation and the Monte Carlo simulation approach. I also discuss different models for volatility estimation, crucial to build appropriate VaR models, ranging from simple moving average methods to more sophisticated GARCH models. Finally, I discuss what is backtesting, why it is adopted to review VaR performance, and what type of tests can be performed.

The second chapter includes a dissertation of the existing literature on Value at Risk relevant for the study that follows. After the presentation of a work that investigates what approaches financial institutions actually adopt, I discuss the results obtained by researchers on the VaR topic. The most investigated aspects are what models perform best for volatility estimation and for VaR estimation.

In the third chapter, I present the empirical study that I have performed in order to assess the performance of the historical simulation approach and the Monte Carlo simulation approach for a portfolio composed of five different asset classes in a stress period for financial markets (from August 2020 to August 2022). I begin by describing the construction of the portfolio, from the data of the securities that compose the portfolio to the weights attached to each instrument; then, I discuss the evaluation framework adopted, in particular how I have divided the data into a rolling window for in-the sample model construction and out-of-the sample for model estimation, and, for the latter, the backtesting that I have performed. I have then proceeded with a detailed analysis of how I estimated the Value at Risk levels under the two simulation approaches. With respect to historical simulation, I cut the relevant daily PnL distribution at the desired percentile (1%), while for the Monte Carlo simulation I have simulated price evolutions assuming different processes for the different asset classes. I then computed the Gross VaR of the portfolio under the two approaches, and the Net VaR of the portfolio under the historical simulation approach. To conclude, I showed the results of the analysis, focusing on how well the VaR models performed.

2 Analysis of Value at Risk

2.1 Risk Management in Financial Institutions

Banks and other financial institutions, when carrying out their activities, bear different types of risks. These risks arise, for example, from deposit taking, loan granting, and security trading. In particular, the most relevant risks are market risk, that is the possibility of losses arising from changes in asset prices; credit risk, defined as the risk of losses due to the borrower's failure to repay the debt or meet contractual obligations; liquidity risk, meaning the capacity to trade assets with short notice and without incurring in great losses; and operational risk, which is the risk of losses arising from the courses of action needed to complete business and financial operations (it may be because of inadequate internal processes, external events, or employee errors) .

Financial risk management is the process of recognizing, measuring and managing the above-mentioned risks. This is of crucial importance for banks in order to ensure that they remain solvent. As a matter of fact, the global financial crisis of 2008 showed how a poor risk containment function can lead to failures of the financial institutions, which generate disastrous consequences for the banking sector, and, in turn, also hurt the real economy in a domino effect. Moreover, the growing interconnection among national financial markets has made it easier to propagate shocks even beyond national borders. Good risk management is thus necessary for financial institutions, in order to ensure that they remain active in providing services to customers, and to avoid spillovers that extend the crises.

With the objective of financial stability, regulators in most areas of the world have thus decided to make sure that banks at all times have layers of capital that work as a buffer in case banks face unexpected losses. The level of capital that banks must hold is commensurate to the level of risks they decide to bear. Specifically, it is linked to the risk-weighted assets (RWAs) of a bank. This means that to each class of the assets of a bank (credit card lending, stocks, bonds, and other types of assets) is assigned a multiplicative factor that represents the risk of loss from that asset for the bank. The greater the riskiness of an asset, the higher the multiplicative factor attached. Once the level of risk-weighted assets is obtained, regulatory capital is a percentage (8%) of RWAs. Regulatory capital consists of three categories: Common Equity Tier 1 (CET 1), Additional Tier 1, Tier 2 Capital. Total regulatory capital is the sum of the three. The distinction among them is based on the loss absorption capacity of the instrument. Capital in the CET 1 is the core capital, the one of highest quality, mainly represented by retained earnings and common shares.

The Basel Committee on Banking Supervision (BCBS)[1] introduced a system of minimum capital requirements, linked to the potential losses caused by the major sources of risk, namely credit risk, market risk, and operational risk. For the fourth major source of risk for financial institution, liquidity risk, the regulator did not impose capital requirements, but a liquidity buffer, that works as a cushion so as to ensure that banks are able to face temporary liquidity shortages. Basel I, the first agreement, actually defined capital requirements only for credit risk, and only with the following agreements (Basel II and Basel III) capital requirements were extended to account also for market and operational risk. A Supervisory Review and Evaluation Process (SREP) has also been established in order to ensure that banks meet the capital requirements and, at the same time, to encourage financial institutions to develop best practices in risk management. While the relevant competent authorities perform Use Tests and Stress Tests to test the stability and resilience of institutions, banks on their part must carry out the Internal Capital Adequacy Assessment Process (ICAAP), in order to autonomously evaluate their capital appropriateness. In case the results of the controls should prove to be not satisfactory, the bank is subject to a capital add-on.

[1]The Basel Committee on Banking Supervision is composed by the chairmen of Germany, Belgium, Canada, France, Italy, Japan, Luxembourg, Netherlands, Spain, Sweden, Switzerland, the United Kingdom and the United States of America. The committee signed a series of agreements, called Basel Accords, providing recommendations and opinions on the intended development of regulation in the banking sector.

2.2 What is Value at Risk

Market risk, which is the focus of this thesis, arises from five main sources: exchange rate risk, interest-rate risk, equity risk, commodity risk, and volatility risk. When the market price of a financial instrument, or portfolio of financial instruments, is affected by one of these variables, we say that the position is sensitive to the performance of these variables.

Assets subject to market risk are those that are placed in the trading book of a bank. Specifically, the securities of a financial institution are allocated either to the trading book or in the banking book. According to the Bank for International Settlements (BIS), banks may only include a financial instrument, instruments on FX or commodity in the trading book when there is no legal impediment against selling or fully hedging it. Otherwise, the instrument must be placed in the banking book. With the proposals contained in the Fundamental Review of the Trading Book (2019), the regulator tried to define a more objective boundary between the trading book and the banking book, thus reducing regulatory arbitrage for banks that opportunistically allocated assets based on convenience. As a matter of fact, the previous boundary depended on the bank's subjective intention to trade the asset.

When the exposure to market risk began to be computed, financial institutions used standardized methods to calculate it. However, they are not sensitive to the bank's risk profile: they only account for nominal values of assets and from these derive the exposure of the bank as a percentage of nominal value.

Nowadays, most financial institutions use the Internal Models Approach (IMA) for market risk calculation. The use of an internal model is subject to the approval by the relevant competent authority under Articles 355 *et seq.* of the Capital Requirements Regulation (CRR). In order to gain approval, a bank must satisfy some requirements: in particular, the institution must develop sound risk management practices and conduct regular stress testing activities, and the internal model must have a proven track record of accuracy. Also, the human capital of the bank must be adequate to ensure that it is able to build internal models that are appropriate. Finally, banks are encouraged to have an independent risk control unit. The most used internal method is Value at Risk (VaR). One of the first institutions to develop a VaR model was JP Morgan in 1995, which produced the *RiskMetrics*™ model[2].

Value at Risk defines what is the maximum loss of a position that will be sustained given a certain confidence level within a defined horizon of time. It has the merit to provide a single number which represents the exposure to market risk in monetary terms. This is beneficial for practitioners because of its ease of use in discussions with top management and reporting with the regulator. In fact, VaR is used by commercial banks, investment banks, pension funds, hedge funds and other financial institutions.

The VaR with confidence level α is given by:

$$VaR(\alpha) = F^{-1}(1-\alpha) \tag{1}$$

where F^{-1} is the inverse of the cumulative distribution function of the returns of a position.

From a probabilistic perspective,

$$Prob(r_t \leq -VaR) = 1 - \alpha \tag{2}$$

where r_t is the return realized by a financial instrument at time t and α represents the confidence level. We generally have $r_t = ln\left(\frac{P_t}{P_{t-1}}\right)$, where P_t and P_{t-1} are, respectively, the price of the instrument at time t and at time $t-1$. The analysis of log-returns is easier than that of prices because returns are mean-reverting while prices are not[3].

Value at Risk has also an important property: the forecast of VaR at time $t+h$ depends

[2]The model developed by JP Morgan assumes that financial returns are distributed normally and are independent and identically distributed. Other models relax this assumption in an attempt to produce better outcomes.

[3]Most financial returns time series are stationary (mean-reversion is one of the characteristics), while price time series are not. The analysis of stationary time series, and their forecasting, is easier. For a stochastic process to be stationary, it must possess three characteristics: constant mean, constant variance, and the correlation between two observation must depend only on the time distance between the two observation, and not the particular time in which they are taken from. In other words, the distribution of a sequence depends only on the relative position of the sequence, not the absolute position.

only on I_t, that is the information set at time t.

Besides the pros of VaR that I have previously mentioned (such as its simplicity of use), other aspects justify the fact that it has gained popularity over standardized approaches to market risk. First, it can be applied to diverse financial instruments, such as stocks, fixed income instruments, derivatives. Also, it takes into account several risk factors at the same time, by providing a measure that accounts for all relevant risk for the specific asset class. When we shift from the VaR of a single instrument to the VaR of a portfolio, also the correlation among the instruments that compose the portfolio must be taken into account.

Nonetheless, Value at Risk also shows some drawbacks: first, it does not tell us anything about what happens when the loss is actually greater than VaR. This may be a serious problem because two portfolios with the same VaR may have a very different left tail behavior. As a consequence, it may be important to further investigate the amount of possible losses beyond VaR. This problem can be solved by using CVaR instead of VaR. CVaR, or Conditional VaR, also called Expected Shortfall, is defined as the average of the loss, given that the loss is greater than VaR. It is thus a conditional expectation. Mathematically, $CVaR = E(r_t | r_t < -VaR)$. CVaR can also solve another disadvantage of VaR: the lack of sub-additivity, which implies that VaR is not compliant with one of the essential features of a consistent risk measure proposed by Artzner et al. (1999)[4]. For a sub-additive risk measure, the risk of a portfolio composed by a number of different positions must not be higher than the sum of the risks of the individual positions.

A risk measure F($\cdot$) is sub-additive if:

$$F(X+Y) \leq F(X) + F(Y) \tag{3}$$

[4]Consistent risk measures have the following characteristics: they are translation invariant (an amount of cash added to a portfolio reduces the risk by the same amount), homogeneous (scaling the size of a portfolio by b scales risk by b), monotonic (if portfolio A has returns that are sistematically lower than the return of portfolio B, its risk must be higher), and sub-additive.

The lack of sub-additivity means that the VaR of a portfolio may be larger than the sum of the VaRs of the instruments that compose the portfolio. For a portfolio of two assets A and B, for example it may happen that $VaR(A+B) \leq VaR(A)+VaR(B)$. Even though for the majority of asset classes, where risk factors are elliptically distributed, such as stocks, sub-additivity is not violated, there are other particular assets for which this is a problem.

2.3 Regulatory Overview

As mentioned in the previous section, regulators have decided to impose levels of capital to cover losses in the balance sheet of banks arising from changes in market prices. For institutions adopting internal approaches for market risk, the level of regulatory capital is linked to the VaR estimate provided by the bank. According to the Basel Committee, VaR shall be calculated under normal market conditions on a daily basis with a 99% one-tailed confidence interval and a 10-day holding period. The Basel Committee then leaves banks free to adopt different approaches for value at risk estimation, the most frequent being variance-covariance, historical simulation approach, and Monte Carlo simulation approach.

The Basel Committee imposed, as minimum capital requirement for banks using VaR, a level of capital not lower than 3 times (multiplicative factor) the VaR. The value at risk measure to be taken into account is the greater value between the previous day VaR or the average of the VaR measures obtained in the last sixty days. The multiplicative factor can be increased to a factor up to 4 in case the outcome of the model proves to be inadequate in predicting losses. The outcome is evaluated through the so-called backtesting: it is the process of comparing the losses predicted by the VaR model to those actually faced by the financial institution over a defined period of time. The Basel Committee has set up the following framework, in which the evaluation is for the last 252 business days and it must be carried out at least on a quarterly basis. If the number of exceptions is below four, the bank is in the green zone, and will not be subject to any additional capital charge. When the number of exceptions is between 5 and 9, the bank is in the yellow zone, and will be subject to an additional charge between 0.40 and 0.85. When exceptions exceed 10, the bank's capital requirement is increased to 4 times the VaR. To be precise, the increase must always be applied if exception incur because of *model integrity*, that is risk positions were reported incorrectly, or *model accuracy*, which means that the model is not able to capture risks in an accurate manner. If instead the exceptions arise because of *intraday trading* (risk positions may change throughout the day), then regulatory authorities should "seriously consider" applying the increase in the multiplicative factor, but this is not mandatory. Table 1 shows in the specific the penalty zones defined by the regulator.

Table 1: The Basel Penalty Zones

Zone	Number of exceptions	Potential increase in k
Green	0 to 4	0.00
Yellow	5	0.40
	6	0.50
	7	0.65
	8	0.75
	9	0.85
Red	$\geq$ 10	1.00

Another aspect worth mentioning is that in 2009 regulators have introduced an incremental risk capital charge (IRC), which covers default risk and migration risk for unsecuritized credit products, and a measure called *stressed* VaR SVaR (and also *stressed* CVaR), to be calculated at least on a weekly basis (institutions must instead calculate VaR daily). The Incremental Risk Charge addresses the risks associated to the downgrade in the rating or default of a security. It is calculated assuming a holding period of 1 year and at a 99.9% confidence level. With respect to Stressed VaR, a bank has to find a 252-day time period in which it experienced extreme stress in its portfolio, for the purpose of SVaR computation. Frequently, a 12-month time frame between 2007/2008 is used to represent such a period of stress in the financial markets. The rationale is that VaR showed a structural problem of procyclicality: tranquil periods may lead to an estimate of VaR that may not be accurate in reflecting the future risk of market losses. The European Banking Authority (EBA) published the Guidelines on Stressed Var (2012) with the objective to enhance the harmonization of supervisory practices among banks.

Banks using internal models may also be subject to an additional, specific risk, capital charge. This represents a buffer against idiosyncratic risk factors, for example in the case the portfolio in consideration is not diversified.

A good estimate of VaR is therefore of crucial importance for banks. On the one hand, if the outcome of a VaR model overestimates the true VaR, the bank will tie up more capital than is actually needed, and this represents a cost for the institution, which foregoes investment opportunities that would be more profitable. On the other hand, if the VaR model underestimates the true VaR there will be two problems: first, the number of exceptions found by the backtesting process will probably put the bank in a yellow or red zone; this will increase the bank's capital requirement and thus, again, an unnecessary cost for the institution. The second, more serious, problem is that if VaR is underestimated, the bank risks to be insolvent if it has put aside a level of capital insufficient to cover losses that may arise.

Financial institutions may thus have incentives to underestimate the VaR to minimize capital requirements (the so-called regulatory arbitrage); given that there are different approaches to compute VaR, the risk is that institutions choose to use the one that benefits them. An harmonization of VaR calculations among different institutions is thus desirable in order to level the playing field among different banks.

With respect to backtesting, despite its merits, it is important to mention the fact that it ignores the time pattern of losses. Specifically, it does not differentiate if exceptions occur evenly distributed or clustered. From a regulatory point of view, exceptions are simply counted over the last year, but no considerations are taken from the distribution of exceptions. However, when large losses are clustered, the risk of insolvency of a bank increases.

As a matter of fact, a series of large losses concentrated in a few days may wipe out the capital of a bank.

The Market Risk Amendment of 1998 (MRA) provided that the 10-day VaR needed for regulatory purposes could be obtained by scaling the 1-day VaR using the square root of time, thus $\sqrt{10}$. It is debatable whether this has a sound theoretical meaning, given that the assumptions that justify such a choice are that risk factors are independently and identically distributed (i.i.d.) across time with zero mean. However, for financial asset classes, this is rarely verified, as returns and volatilities tend to cluster.

With respect to Value at Risk computation with portfolios composed by options (relevant in the following chapters), the Basel Committee on Banking Supervision prescribes banks to capture the non-linear characteristics of option prices and build models that capture vega risk, meaning the level of risk of options that arises by changes in the volatility.

2.4 Standard Approaches to VaR computation

In this section, I analyze the three most common methods used by financial institutions, namely the Variance-Covariance Approach, Historical Simulation, and Monte Carlo Simulation. These methods are called standard methods and are the ones that most banks actually use, and the first to be developed, even though the computation of VaR can take other, more complex forms, that will be further analyzed in the following chapters, with an empirical study of VaR performance under different models.

2.4.1 Variance-Covariance Approach

Under the Variance-Covariance Approach Value at Risk is estimated through a closed form formula.

The simplest form of the Variance-Covariance method, the *delta-normal* approach, is applicable to portfolios composed by assets whose return is a linear function of the returns of risk factors or underlying assets.

The formula for a single stock position is:

$$VaR = MV \cdot k \cdot \sigma \cdot \sqrt{t} \tag{4}$$

where MV is the current market value of the position, k is the multiplicative factor related to the number of standard deviation according to α, σ is the annual volatility, that is scaled by the square root of time $\sqrt{t}$. For the regulatory purposes of $\alpha = 99\%$, under the normal distribution assumption, k equals 2.33.
In more general terms, if the asset is not a stock, we may add a δ in the formula to account for the sensitivity of the position's market value to changes in the market factor. The formula will thus be: $VaR = MV \cdot \delta \cdot k \cdot \sigma \cdot \sqrt{t}$.

When we shift from a single position to a portfolio, the correlation among the different instruments must be taken into account. VaR_p of a portfolio is thus:

$$VaR_p = \sqrt{\sum_{i=1}^{N}\sum_{j=1}^{N} VaR_i VaR_j \rho_{i,j}} \tag{5}$$

where $\rho_{i,j}$ is the correlation coefficient between two financial instruments i and j. It is important to note that since $\rho_{i,j} \leq 1$, $VaR_p \leq \sum_{i=1}^{N} VaR_i$. VaR computed using the variance-covariance approach is thus sub-additive.

When we analyze VaR for portfolios, it is easier to work with matrices. The VaR of N different positions is the vector:

$$v = \begin{bmatrix} VaR_1 \\ VaR_2 \\ ... \\ VaR_N \end{bmatrix}$$

The correlation coefficients can be shown in the matrix:

$$C = \begin{bmatrix} 1 & \rho_{1,2} & ... & \rho_{1,N} \\ \rho_{2,1} & 1 & ... & \rho_{2,N} \\ ... & ... & ... & ... \\ \rho_{N,1} & ... & ... & 1 \end{bmatrix}$$

Thus, eq. (37) can be rewritten as:

$$VaR_p = \sqrt{v^\top \cdot C \cdot v} \tag{6}$$

where $v^\top$ is the transpose matrix of v.

For some instruments, the relationship between the shocks in the market factor and changes in the market value of the instrument is not linear. To consider this aspect, we must add a second-degree or *curving* term. A bond's price, for example, does not vary linearly as the yield to maturity changes; for this reason, we also take into account convexity. Similarly, for options, besides the above-mentioned δ coefficient, we have to introduce in the analysis the γ coefficient. This approach is called the *delta-gamma* approach.
In a *delta-normal* approach,

$$\Delta MV \simeq \delta \cdot \Delta R \tag{7}$$

where ΔMV is the change in the market value of the instrument and ΔR represents the change in the risk factor.
In a *delta-gamma* approach,

$$\Delta MV \simeq \delta \cdot \Delta R + \frac{\gamma}{2} \cdot (\Delta R)^2 \tag{8}$$

Adding the gamma coefficient improves the approximation; this is particularly true for instruments with a high γ, such as at the money options approaching maturity.

The analysis presented up to here assumed that positions were sensitive to only one market factor. In reality, the price of an asset may be sensitive to several market factors. In this case, the position is broken down into different components in order to study the individual effect of each elementary component. The impact of each component is then aggregated to end up with a measure of exposure, also taking into account the correlations among the market risk factors. This technique is called mapping of risk positions. To make an example, an Italian bank which has in its balance sheet a US ten-year Treasury bond is exposed to two different sources of risk: the EUR/USD exchange rate and the level of dollar interest rates. Here it is necessary to compute the VaRs for two positions, namely the change in value on a foreign currency position, and the change in value on a foreign bond held by a foreign investor, so as to eliminate exchange rate risk. Once the two VaRs are obtained it is possible to compute the VaR of the position, by summing the two and then multiplying the result for the correlation coefficient between the two market factors.

The main reason why the variance-covariance approach has gained popularity is its simplicity in terms of computation: other approaches are complex because of their calculation intensiveness. However, this model has some drawbacks: first, the assumption of normally distributed returns is in contrast with the empirical behavior of many financial assets, including stocks. As a matter of fact, financial returns often show negative skewness (which means that the distribution is skewed to the left) and significantly excessive kurtosis (which means fat tails). Therefore, the size of the actual losses may be larger than what predicted by the normal distribution. To account for fat tails a t-Student distribution is sometimes used instead of the normal. However, it is important to note that even if the returns of an asset were not normal, the returns of a well-diversified portfolio would anyway be normally distributed if risk factors were independent[5], even though the returns of each security are not. Nonetheless, also this assumption is implausible. Similarly, the

[5]This result is due to the Central Limit Theorem, which states that the sum of i.i.d. random variables converges in probability to a standard normal N(0,1), even if the random variables were not normally distributed themselves.

assumption that variance is constant contradicts the actual behavior of most financial time series.

2.4.2 Historical Simulation

In the historical simulation approach, future price changes are assumed to be properly represented by their empirical historical distribution. The rationale is that past behavior reflects the future behavior of market variables. Specifically, the procedure to compute VaR using the historical simulation method is the following:

1. Choose a sample of return for a security or a portfolio of securities over a given period (it is necessary to identify appropriate intervals, for instance, daily, monthly, etc. and a proper time-frame).

2. Revalue the security or the portfolio at each observation.

3. Sort the returns from the lowest to the highest and "cut-off" the distribution at the percentile corresponding to the confidence level needed.

4. The VaR is thus given by the difference between the value found in the previous step and the current security or portfolio market value.

The main pros of this approach are the following: first of all, it is very easy to implement. Also, it can accommodate for fat tails, skewness and other characteristics that deviate from the normal distribution, if these characteristics are present in the historical data set. Differently from the variance-covariance approach, it is not needed to assume the distribution of returns, since it is retrieved from the historical one. Moreover, it does not require to estimate the variance-covariance matrix. Being based on full valuation of the position, the output is not an approximation such as the delta-normal approach. Moreover, it is the easiest to understand, which is beneficial also in internal reports. For instance, an increase in VaR obtained using this approach is more easily communicated to top management.

The limitations of this approach are instead the following: first, its results are completely dependent on the data set. For instance, it is appropriate only as far as the returns of the asset are stable over time; in particular, if the distribution of returns is heteroskedastic (i.e. shows a time-changing variance), then the empirical probability distribution will fail to capture adequately future market movements. A critical issue when using the historical simulation approach is the time horizon to consider: first, it is important to choose a proper time span. In fact, if it is chosen too short, it may not be wide enough to contain relevant possible realizations of market movements; if instead it is too long, it may include realizations that are not indicative of future possible evolutions (they may be obsolete). Moreover, for some time series there may not be enough data to make the analysis sound from an econometric perspective, in particular if the frequency of observations is not daily. In this case the results obtained will not be statistically significant. A practical drawback of this approach is that it only provides VaR estimations at discrete confidence intervals, thus it may be impossible to have some confidence levels in case the number of observations is not high enough. It is important to note that this approach treats equally all the realizations without considering that newer ones may be more indicative for predicting future behavior. To conclude, it is quite intensive in terms of computation, and it may take a long time to estimate VaR.

2.4.3 Monte Carlo Simulation

The last standard approach for Value at Risk measurement that I am going to analyze in this chapter is the Monte Carlo simulation. It consists in generating random data given a defined probability density function (the one that best fits the data of returns), and, using the same percentile logic as the historical simulation, rearrange observations and cut them at the desired percentile to find the VaR. A large number of simulations is needed in order to minimize sampling variability due to randomization. As a matter of fact, the

Law of Large Numbers[6] ensures convergence to a stable VaR output when the number of iterations increase to a very high number.

Specifically, the steps to follow to obtain the VaR of a single instrument are the following:

1. Choose the probability distribution that best fits the distribution of returns.
2. Estimate the parameters of the distribution that are needed to simulate scenarios; for, example to simulate a normal distribution the parameters to estimate are the mean and the standard deviation.
3. Simulate N scenarios from the selected distribution; generally N equals 10,000.
4. Calculate the change in the position's market value for each simulated scenario.
5. Cut the distribution at the percentile corresponding to the confidence level selected.

When we move from a single instrument to a portfolio, we must select a joint probability density function, in order to account for the different market factors that influence the market value of the portfolio. Also, for portfolios value at risk it is necessary to calculate the covariances among the market factors before simulating the scenarios.

Generally, it is assumed that the price of some assets, in particular stocks, follow a geometric Brownian motion process.

$$\frac{dS}{S} = \mu dt + \sigma dz \tag{9}$$

where S is the price of the stock, dS is the instantaneous change in the price of the stock, µ is the expected annual return of the stock, dt is the infinitesimal interval of time, and σdz is the noise, or stochastic component, for the unpredictable future evolution of the stock price. The shocks are assumed to be independent and identically distributed. Because of Ito's lemma, returns $\frac{dS}{S}$ are normally distributed.

[6]The Law of Large Numbers is the theorem that guarantees that as the number of identically distributed, randomly generated variables increases, their sample average approaches the true population mean.

A first merit of this approach is the full valuation of positions, as it happens with historical simulation. This makes it a suitable approach for non-linear portfolios, because it does not require simplifying assumptions about the joint distribution of the underlying data. Also, differently from the variance-covariance approach, that is based on the normal distribution, Monte Carlo simulation can be used for any probability distribution. This leaves flexibility in selecting the correct probability density function.

One of the main disadvantages of this approach is that it is time consuming and computationally intensive. Estimating a large number of scenarios may take a long time, in particular for positions such as exotic options. However, it must be noted that for such positions, using a *delta-normal* approach, would provide an unsatisfactory VaR estimate. Another limitation is that, contrary to the historical simulation, it requires the estimation of a variance-covariance matrix.

2.5 Volatility Estimation Models

When adopting the Variance-Covariance approach and the Monte Carlo simulation approach, an element of crucial importance is the estimation of volatility. In fact, we need volatility as an input both to use the formula that gives us the VaR under the Variance-Covariance approach, as well as to have the parameters needed to generate random scenarios under the Monte Carlo simulation approach.

There are basically three methods to estimate volatility. The first is to use historical volatility, thus assuming that volatility is a constant parameter. However, studies show the volatility of returns of financial variables is not constant. The second possibility is to construct a more sophisticated model that allows volatility to change over time. Past volatility here is used to construct the volatility estimation model, but does not coincide with it. Different classes of models exist, from simple moving average (MA) processes to exponentially weighted moving average (EWMA) processes, to generalized autoregressive conditional heteroskedasticity (GARCH) processes. The third method consists in using the volatility implied in option prices. In this case, instead of using a model with volatility as an input to find the option price, the option price is used as an input to find the implied volatility that justifies such a price. The merit of this last approach is that it gives a forward-looking measure of variance, whereas the other methods provide a backward-looking variance, since they are constructed based on historical data. However, the practical drawback of the implied volatility approach is that it can only be used as long as there exists an option contract whose underlying asset is the same as the one for which volatility is intended to be forecasted, and this option contract is traded on a liquid, efficient and organized market. For this reason, this approach is in practice rarely used for volatility estimation.

In the following section I am going to analyze the moving averages and GARCH models, which are more relevant for the empirical study that I will present in the following chapters.

2.5.1 Moving Average Models

The use of moving average models to estimate volatility is the most widespread method by practitioners. Volatility is measured as the square root of variance using a sample of n observations.

$$\sigma_t = \sqrt{\frac{\sum_{i=t-n}^{t-1} (r_i - \overline{r}_t)^2}{n-1}} \tag{10}$$

where $\overline{r}_t$ is the sample mean of the position calculated at time t using n observations from time $t-n$ to time $t-1$. The model is called a moving average because as we move to the following period estimation of σ, we will drop the oldest observation and replace it with the last, thus keeping fixed the number of n observations, moving the sample's time window forward by one period.

The pitfall of this method is that the weight assigned to the n observation is the same for all n; however, a good model shall take into account the fact that more recent observations are more likely to be indicative of future variance. This problem is solved by the Exponentially Weighted Moving Average (EWMA) models, which attaches greater weights to the most recent observations. The estimate of volatility is thus made more responsive to recent shocks. Analitically, volatility is:

$$\sigma_t = \sqrt{\frac{1-\lambda}{1-\lambda^n}\sum_{i=0}^{n-1}\lambda^i r_{t-1-i}^2} \tag{11}$$

where λ is a decay factor between 0 and 1 which relates to the persistence of a past shock in today's volatility. A lower λ produces a VaR output that is more affected by recent conditions, and, thus, more volatile. If lambda is sufficiently small and/or n large, then $\lambda^n \simeq 0$. In this case, the volatility can be approximated to:

$$\sigma_t = \sqrt{(1-\lambda)\sum_{i=0}^{n-1}\lambda^i r_{t-1-i}^2} \tag{12}$$

The practical issue with this approach is to estimate λ and to select the number of past observations to include in the model. The previously mentioned *RiskMetrics*™ model developed by JP Morgan uses $\lambda = 0.94$ and $n = 74$.

2.5.2 GARCH Models

Another possibility to estimate volatility is to use GARCH models, which have the merit to accommodate for conditional heteroskedasticity, namely the fact that periods of high volatility follow quite periods of low volatility and vice versa (this path-dependent behavior is called volatility clustering).

The GARCH model stands for Generalized Autoregressive Conditional Heteroskedasticity. Generalized refers to the fact that GARCH is a generalization, proposed by Bollersev in 1986, of the ARCH model devised by Engle in 1982. The merit with Bollersev's formulation is that it is a more parsimonious model: it requires an estimation of less parameters and it is thus more stable over time. Autoregressive means that a regression is made on the dependent variable itself (for our purpose the variance). Conditional heteroskedasticity refers to the fact that variance is not constant, and, in particular, the estimates are obtained based on the information set available in the previous period.

The analytical formulation of a GARCH (p,q) process is as follows:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2 \tag{13}$$

where ϵ_i represents the prediction errors and σ_t^2 is the estimate of the variance. Previous periods shocks and previous periods variance are lagged, respectively, by i and j periods. The prediction errors ϵ_i are assumed to be normally and independent and identically distributed and they have the following characteristics:

$$E[\epsilon_t | I_{t-1}] = 0 \tag{14}$$

$$VaR[\epsilon_t | I_{t-1}] = \sigma_t^2 \tag{15}$$

where I_{t-1} is the information set available at time $t-1$. GARCH model thus keep the normality assumption but not for the market returns; instead, the hypothesis is that the prediction errors are normally distributed.

The parameters of GARCH models are generally estimated through the Maximum Likelihood Estimator (MLE). This is accomplished by maximizing a probability function such that the observed data is the most likely under the assumed statistical model. The majority of applications of the GARCH model is based on the GARCH (1,1). For financial variables, we have:

$$r_t = \epsilon_t \tag{16}$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \tag{17}$$

What the model means is that returns, especially in a short period (say one day) are unpredictable with zero mean, and then a GARCH model is used in order to estimate volatility. It is important to mention that, even if returns followed a more complex process (for instance, they may be serially correlated), it is possible to standardize returns by dividing them by the estimate of the conditional volatility. Analitically, $e_t = \epsilon_t / \sigma_t$. In this case, the standardized returns are independent and identically distributed.

The GARCH model assumes that volatility is stationary and there exists a long-run variance to which it reverts towards. In particular, it is possible to show that long term unconditional variance is given by:

$$\sigma^2 = \frac{\alpha_0}{1 - \alpha_1 - \beta_1} \tag{18}$$

It is thus possible to rewrite the GARCH (1,1) model as:

$$\sigma_t^2 = (1 - \alpha_1 - \beta_1)\sigma^2 + \alpha_1 \epsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \tag{19}$$

This type of representation is convenient because it decomposes the estimation of the conditional variance into a weighted average of long-run variance, the expectation of the variance for the previous period, and a non-predictable shock for the previous period.

Another way of reformulating (17) highlights another important feature of GARCH models: that they present an infinite memory, even though they rely on only one lag of the variables in the model. It is possible to show this in two passages:

$$\sigma_t^2 = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta_1(\alpha_0 + \alpha_1 \epsilon_{t-2}^2 + \beta_1 \sigma_{t-2}^2) \tag{20}$$

By means of a recursive substitution we obtain:

$$\sigma_t^2 = \frac{\alpha_0}{(1 - \beta_1)} + \alpha_1 \sum_{i=1}^{\infty} \beta_1^{i-1} \epsilon_{t-i}^2 \tag{21}$$

By equation (21) it is evident that a shock, no matter how far in time, still affects the estimate of variance.

One pitfall of this model is that it does not make a distinction between positive or negative shocks. However, what banks are in practice interested into are negative shocks. Also, empirical studies find that volatility increases at a higher rate when returns are negative compared with when they are positive. To take this asymmetry into account, several extention of the GARCH model have been developed, such as Exponential GARCH (EGARCH). Usually a EGARCH (1,1) model is used, whose formulation is:

$$ln(\sigma_t^2) = \alpha_0 + \gamma\Big(\frac{\epsilon_{t-1}}{\sigma_{t-1}}\Big) + \alpha_1\left(\left|\frac{\epsilon_{t-1}}{\sigma_{t-1}}\right| - \sqrt{\frac{2}{\pi}}\right) + \beta_1 ln(\sigma_{t-1}^2) \tag{22}$$

where γ is the so-called leverage parameter. A negative value of γ implies that previous negative shocks have a greater impact on current conditional volatility than previous positive shocks. Therefore, for financial time series applications we expect $\gamma < 0$. For the model to be properly specified, the restriction $\alpha_1 + \beta_1 < 1$ must apply.

2.6 Backtesting

As I have already mentioned in the previous section, the Basel Committee requires that banks using internal models for market risk backtest the model adopted to check its validity. In particular, they should ensure that the model's predicted losses are consistent with the losses actually realized. The standard tests imposed by the regulator evaluate the accuracy of the model, and there are mainly two types of tests: the Unconditional Coverage Test and the Conditional Coverage Test. Backtesting tests may be used also to compare loss functions among different models. However, the Basel Committee is only interested in the first approach for the purpose determining capital requirements.

The Unconditional Coverage Test was developed by Kupiec in 1995. The null hypothesis to be tested is that the frequency of empirical exceptions is consistent with the confidence level adopted, which reflects the number of "theoretical" exceptions α. The test is called unconditional because it does not take into account the sequence of occurrence of exceptions across time (i.e. the probability of an exception is not evaluated conditional to an exception in the previous period). This is a one-tailed test with the alternative hypothesis represented by an empirical exception rate greater than α. To implement the test we must define an indicator function[7] of exceptions:

$$I_{t+1} = \begin{cases} 1 & \text{if } r_{t+1} < \text{VaR} \\ 0 & \text{if } r_{t+1} \geq \text{VaR} \end{cases} \tag{23}$$

Assuming a constant probability of exceptions, the number of exceptions $x = \sum I_{t+1}$ follows a Binomial distribution B(N,α) where N is the number of observations in the sample. A model is considered accurate if the null hypothesis $\hat{\alpha} = \alpha$ where $\hat{\alpha} = \sum I_{t+1}/N$ is not rejected. This can be estimated through a likelihood ratio test. The statistic LR_{uc} is thus obtained:

$$LR_{uc} = -2ln\left[\frac{\alpha^x(1-\alpha)^{N-x}}{\hat{\alpha}^x(1-\hat{\alpha})^{N-x}}\right] \tag{24}$$

which follows a $\chi^2(1)$ distribution that is a chi-square distribution with 1 degree of freedom.

In 1998 Christoffersen developed a Conditional Coverage Test, in order to solve the problem of the Unconditional Coverage Test, namely the fact that it does not take into account the time distribution of exceptions. Christoffersen introduced a test for serial dependence of exceptions. The Conditional Coverage Test jointly test whether the average number of exceptions is adequate and whether they are serially independent. The

[7]The indicator function of an event is a random variable that takes value 1 when the event happens and value 0 when the event does not happen.

likelihood ratio is given by:

$$LR_{ind} = -2ln\left[\frac{L(x|\hat{\alpha})}{L(x|\hat{\alpha}_{0,0}, \hat{\alpha}_{1,0}, \hat{\alpha}_{0,1}, \hat{\alpha}_{1,1})}\right] \tag{25}$$

where:

$$L(x|\hat{\alpha}) = (1 - \hat{\alpha})^{N-x}\hat{\alpha}^{x} \tag{26}$$

The test statistic LR_{cc} is given by $LR_{cc} = LR_{uc} + LR_{ind}$ and it is distributed as $\chi^2(2)$. Specifically:

$$LR_{cc} = -2ln\left[\frac{L(x|\hat{\alpha} = \alpha)}{L(\hat{\alpha}_{0,0}, \hat{\alpha}_{1,0}, \hat{\alpha}_{0,1}, \hat{\alpha}_{1,1})}\right] \tag{27}$$

Alternatively, a *backtesting criterion* statistic Z can be used:

$$Z = \frac{N\hat{\alpha} - N\alpha}{\sqrt{N\alpha(1-\alpha)}} \tag{28}$$

which is distributed as a standard normal N(0,1).

For risk management purposes, financial institutions may also be interested in measuring the amount of losses incurred in case returns are lower than VaR. In fact, when an exception occurs, the level of the loss is relevant for the bank. Following this rationale, the best VaR model is the one that minimizes the loss function. The first type of approach presented does not give an indication with this respect. Thus Lopez (1999) suggested a test which measure the loss function L_{t+1} in the following way:

$$L_{t+1} = \begin{cases} 1 + (r_{t+1} - VaR)^2 & \text{if } r_{t+1} < \text{VaR} \\ 0 & \text{if } r_{t+1} \geq \text{VaR} \end{cases} \tag{29}$$

If different VaR models are compared, the one with the lowest loss function over a certain time period is preferred.

3 Previous Literature on VaR

In this section, I am going to present a brief summary of the vast literature that concerns Value at Risk. I will mainly analyze the methodolgy (data and time period used, VaR models investigated, etc.) and the results of empirical studies that focus on the performance of different VaR models. Although the numerous researches carried out in the past years, no consensus exists among researchers or institutions with respect to the best VaR approach. Among the number of empirical studies carried out by researchers on this topic, I will predominantly discuss those that are more relevant for the purpose of the thesis.

To begin with, it is interesting to examine a paper by Pérignon and R. Smith (2010) in which the authors analyze the level and quality of Value at Risk disclosure by commercial banks. They build a VaR Disclosure Index (VaRDI), which yields a summary on the quantity of VaR disclosure made by banks and the level of information provided by them to enable interpretation of VaR, including information regarding VaR construction. Specifically, VaRDI is composed by six factors: VaR basic characteristics (holding period and confidence level), summary VaR statistics (high, low, average, year-end VaR, VaR by risk category), summary information about the previous year's VaR, histogram or plot of daily VaRs, and backtesting (number of exceptions). The researchers consider a cross-section of 60 US, Canadian and international banks and find a great level of heterogeneity across different nations (the most striking difference is that Chinese banks did not disclose VaR figures at all). The most widely used method in the sample was found to be the historical simulation approach, with 73% of the banks adopting this method; the authors also show that this approach poorly estimates future volatility. This has to be kept in mind when results on historical simulation-based VaR estimates will be presented. Also, the authors find an upward trend in the quantity of information released by commercial banks to the public over the decade 1996 to 2005. However, the quality of information disclosed, as represented by the VaRDI index, did not show significant improvements.

The first studies on the performance of historical simulation approach, namely Beder (1995), Hendricks (1996), and Pritsker (1997), found out that this approach performed as good as the other methods confronted, the variance-covariance and the Monte Carlo simulation. Specifically, Hendricks (1996) applied different VaR models to 1000 randomly chosen spot foreign exchange portfolios throughout the period 1983-1994. Then, he compared the performance of three approaches: historical simulation, equally weighted moving average approaches, and exponentially weighted moving average approaches. The author analyzed the accuracy of different models following these three approaches, for example by constructing VaR measures over different time spans, namely 50, 125, 250, 500 and 1250 days. Of course, the VaR estimates are found to be much more volatile when the time spans are shorter, which makes VaR prone to rapid swings that can cause exceptions more easily. The conclusions they reach is that neither of the three approaches outperforms the others. Another finding is related to the level of losses incurred by institutions: while almost all of the methods yield accurate 95th percentile risk estimates, while the 99th percentile risk assessments are less trustworthy, covering just 98.2 percent to 98.5 percent of the events on average. The authors also estimate that, in order to cover precisely 99% of the realizations, risk measures would need to be augmented by 10 percent or more, which is difficult on a level-of-capital basis. To make the scenario even more scary, when an exception arises, the level of capital needed to cover the loss compared to the VaR estimate is 30 to 40 percent larger. However, it is important to mention the fact that these early studies on the performance of different Value at Risk approaches could not include models developed later that may have been able to provide more accurate risk estimates (and thus outperform the models that the authors test).

Raaji and Raunig (1998) perform a comparison of six different Value at Risk approaches and compare the results by constructing 20 different portfolios composed by 100 billion dollars in thirteen foreign currencies, analyzing exchange rates in the period between 1986 and 1998. The first result they obtain is that, on average, the historical simulation with 1,250 days of historical data produces the highest VaR. Also, they find that differences in VaR levels under different approaches sometimes exceeded 200% when the methods are compared with the EWMA-based variance-covariance approach as benchmark[8].

More recently, other studies report that Value at Risk outputs produced through historical simulation are not satisfactory. Such studies were carried out, among the others, by Ashley and Randal (2009), Trenca (2009), Angelidis et al. (2007) and Abad and Benito (2013). The latter study reported that if filtered historical simulation is adopted, which means that volatility is treated as time-varying and an approach to model it conditionally to the information set available is adopted, then the VaR estimate obtained is much more precise than the one obtained with simple historical simulation. The authors employ daily closing prices of eight indexes, namely Spanish IBEX35, French CAC40, German DAX, UK FTSE100, US Dow Jones Industrial Average (DJAI), S&P 500, Japanese Nikkei 225 and Hong Kong Hang Seng (HSI) for the time period from 1994 to 2011. They produce VaR estimates with different models with an innovative feature: the analysis is carried out by taking into account both a stable period (between 2004 and 2006) and a volatile period (between 2007 and 2009). A good VaR model must be *time-robust*, which means that it must pass the backtesting also in stress times. Specifically, they consider a VaR method accurate if and only if it is the accuracy of the model is not statistically rejected in neither of the two time spans under analysis. The first, expected, conclusion from their study is that VaR estimates are more precise during stable periods than during volatile periods. They also find that among the different models evaluated, the Parametric-t-EGARCH[9] is the one that minimizes the loss function, thus supporting the idea that adopting stochastic volatility models for the estimation of volatility improves the VaR estimate.

Pritsker (2006) showed that adopting historical simulation VaR without including time-varying volatility might dangerously under-estimate risk. As a matter of fact, since historical simulation only accounts for the unconditional distribution of risk factors, it is under-responsive to changes in conditional risk. Volatility prediction for most equity, fixed income, and foreign exchange assets decays rapidly with time horizon, according to

[8] The EWMA variance-covariance approach with $\lambda = 0.94$ was used as benchmark because of the wide application of the Riskmetrics model.

[9] By Parametric-t-EGARCH the authors mean a parametric approach to VaR where volatility is estimated trough the means of a GARCH model and errors are assumed to be t-distributed.

Christoffersen and Diebold (2000). The result is that when the VaR horizon is long, capturing time-varying volatility may be less significant than when the VaR horizon is short.

Zikovic and Aktan (2009), Angelidis et al. (2007), Kuester et al. (2006), and Marimoutou et al. (2009) state that the filtered historical simulation (FHS) approach yields the best results among all the methods that they employ. In particular, Zikovic and Aktan (2009) analyze daily log returns series for two currency indexes, namely Turkish XU 100 and Croatian CROBEX for the period 2000-2008. The models under analysis are the following: simple variance-covariance approach, *RiskMetrics* model, historical simulation with rolling windows of 250 and 500 days, BRW[10] (time weighted) simulation assuming decay factors of 0.97 and 0.99, *RiskMetrics* model augmented with GARCH-type volatility forecasting, unconditional Extreme Value Theory (EVT)[11] approach using Generalized Pareto distribution (GPD), conditional quantile EVT approach and, to conclude, what the authors call Hybrid Historical simulation, that is, in fact, a filtered historical simulation (FHS). The conclusions they reach is that only filtered historical simulation and EVT provide acceptable estimates, while all other models would fail the Basel regulatory tests. However, FHS satisfies the backtesting criteria at a much lower cost of capital than EVT models. They also find that the other models not only failed to pass the Kupiec test on the number of exceptions, but they also missed the independence test, meaning that exceptions were clustered and not independent and identically distributed, which leads to models that prove to be too weak to be used in turbulent times.

Barone-Adesi et al. (1999) further proposed a position-by-position FHS method and suggested filtering each risk factor individually and generating volatility predictions for each one. EWMA and the weights proposed by Boudoukh et al. (1998) can also be used in the same way. When weighting or filtering risk variables independently, the implicit

[10]This type of simulation takes the name by the initials of the authors of a paper on this, namely Boudoukh, Richardson, and Whitelaw and, differently from simple historical simulation, attempts to capture time-varying volatility.

[11]The Extreme Value Theory (EVT) method is concerned with the limiting distribution of extreme returns that have been observed over a long period of time and is independent of the distribution of the returns themselves. Since VaR focuses on tail behavior, it can estimate risk more efficiently.

assumption is that the correlation structure across risk factors remains constant over time.

However, evidence of the superiority of FHS is ambiguous; in fact, Nozari et al. (2010) find out that filtered historical simulation does not produce more accurate VaR estimates than other approaches.

Studies conducted on the performance of the variance-covariance approach find the flaws that it carries on a theoretical basis. In fact, as said in the previous chapter, there are some drawbacks connected to this method that impact its reliability. First of all, the fact that it is based on the assumption of normality of returns. However, most financial returns display negative skewness and excess kurtosis which are statistically significant and should not be overlooked if one intends to build a robust model. Black (1976), Pagan and Schwert (1990) criticize the *RiskMetrics*™ model for not taking into account the asymmetry of returns and the leverage effect. Pagan and Schwert (1990) carry out a research studying different volatility models. Taking the 1835-1925 period as the sample on which volatility is analyzed, it emerged that the non-parametric procedures tended to give a better explanation of the squared returns than any of the parametric models. Moreover, the variance-covariance approach imposes i.i.d. returns, which has been proved false in several studies, such as Hansen (1994), Harvey and Siddique (1999), Jondeau and Rockinger (2001). After having observed that returns are not i.i.d., Jondeau and Rockinger (2001) develop a method to measure the conditional dependency of financial returns. This is done through copula functions[12] chosen to follow GARCH-type models, with the peculiarity that they allow for time-varying second, third and fourth moments. The authors take data for American, European and Asian companies; for European companies, the results show not only persistence in dependency but also that large returns of either sign lead to higher dependency in the following periods.

Another aspect investigated in the literature is whether adopting a distribution of returns

[12] Copula functions connect the marginal distributions to their joint distributions and are useful in simulating the linear or nonlinear relationships among multivariate data. A copula is a multivariate distribution function with marginally uniform random variables on [0, 1].

different from the normal distribution improves the VaR estimate. The evidence of the studies does not provide a clear answer. As a matter of fact, some studies, such as Abad and Benito (2013), Polanski and Stoja (2010), Angelidis et al. (2007), report an enhancement in the results when assuming return distributions with tails fatter than those that characterize the normal distribution (such as the t-distribution), while other studies carried out, among the others, by Guermat and Harris (2002) and Billio and Pellizon (2000) conclude that the t-distribution performs equally or worse than the normal distribution. Angelidis et al. (2007) analyze a number of different volatility models by examining their ability to forecast Value-at-Risk (VaR); VaR is estimated for large and small capitalization stocks using non-parametric, semi-parametric, and parametric approaches, and is modeled for long and short trading positions using the Dow Jones (DJ) Euro Stoxx index. Models are compared over two different time periods (stable and turbulent) in order to investigate whether the risk management techniques are robust over time. The first empirical finding they obtain is that the variance-covariance approach does not pass the backtesting in both periods. Moreover, they find that ARCH models under the t-Student distribution overestimate VaR in the majority of cases. Conversely, the filtered historical simulation approach combined with the ARCH volatility specifications provides a significant improvement over either the parametric or the non-parametric methods, and is thus the best model to adopt according to the study.

In fact, when GARCH-type models are adopted for volatility estimation, the empirical evidence seems to suggest that VaR estimates improve. EWMA models prove to perform poorly in several studies, such as Abad and Benito (2013), Ñiguez (2008), Alonso and Arcos (2006), Gonzàlez-Rivera et al. (2004). The VaR output is enhanced even more when asymmetric GARCH models are used as testified by Bali and Theodiossou (2007).

Gurmat and Harris (2002) proposed an exponentially weighted likelihood model, as they pointed out that, for three equity portfolios (U.S., U.K. and Japan), it calculated the VaR more accurately than that of the GARCH model under either the normal or the t-student distributions. Bali and Theodossiou (2006) combined the skewed generalized t-student distribution with ten different GARCH specifications and argued that the t-

GARCH, proposed by Taylor (1986) and Schwert (1989), and the EGARCH had the best overall performance, accurately estimating both VaR and the Expected Shortfall.

Another interesting study has been carried out by Berkovitz and O'Brien (2002); in their research the authors perform a comparison between some internal VaR models actually developed and used by banks with a parametric GARCH model estimated assuming normality of returns. The results they obtain is that banks internal models are not better. As a matter of fact, the GARCH model generally provides for lower VaRs and is better at predicting changes in volatility. Because of this, the GARCH model permits comparable risk coverage with less regulatory capital.

To estimate time-varying volatilities and correlations, multivariate GARCH models have also been proposed, such as the BEKK model proposed by Engle and Kroner (1995) or the DCC model developed by Engle (2002). However, it is important to remark that when there are a large number of risk factors the validity of such models is difficult to analyze. For example, Engle and Sheppard (2007) suggest to average likelihoods before the estimation of the GARCH model through the maximum likelihood estimator (MLE). Also, Engle and Kelly (2009) apply a constraint to the correlation structure that makes the estimation of the model easier while still allowing correlations to vary over time. Furthermore, Aramonte, Rodriguez and Wu (2013) compute Value at Risk for large portfolios by using dynamic factor models to reduce the dimension of risk components and then they estimate a time-varying volatility model. The result that they obtain is that their VaR estimate proves to be superior to both historical simulation and filtered historical simulation. However, rather than adopting even more complex solutions to solve this shortcoming, the banking industry seems to prefer the adoption of simpler solutions, such as equally weighting of observations or reducing the time frame of the data set. Even though these methodology is less accurate, it is much more computationally efficient, particularly for large and complex portfolios.

Hao Li, Xiao Fan, Yu Li, Yue Zhou, Ze Jin, Zhao Liu (2011) carried out a comparison of several different approaches to VaR. Specifically, they constructed a portfolio in the following way: starting from a stock and an option on the stock, they obtained 41 of such pairs. This way they obtained a portfolio well-diversified in terms of industries; not in terms of region, since all the stocks were picked among actively-traded in the US stocks. At-the-money options were used to construct the portfolio. The sample of data spans from 2005 to 2012, but for the purpose of testing the goodness of VaR models, only the second half of the data set was used. First of all, they tested the performance of *Baseline Models*, with the historical simulation approach and the Monte Carlo simulation approach. The outcomes produced in this way show levels of exceptions close to 10%, which makes it evident that both models failed the test. Then, the researchers constructed more sophisticated models and compared their relative performance. When applying filtered methods (bootstrap, empirical and Maximum Likelihood Estimator), under the assumption of t-distributed returns, the outcome is a much better result when compared to normally-distributed returns. Among the three, the MLE method is the one which provides the best forecast. The study also compared different GARCH models for volatility estimation. In particular, GARCH, GJR GARCH and EGARCH volatility models were constructed under the t-distribution assumption; the result is that no model outperforms the others. To sum up, the finding in this study is that filtering historical returns produces the best VaR estimate, which is enhanced further by the t-distribution assumption.

A study carried out by Ribja, Tsagris and Mhalla (2016) finds different evidence with respect to the results provided by the use of different GARCH models. They take a portfolio consisting of five shares, with equal weights, constant over time and they use a 6-year period consisting of 1500 observations. The conclusion is that taking into account the leverage effect, thus building asymmetric GARCH models, actually makes the estimation more accurate. Although not all studies confirm this, evidence seems to be in favor of this conclusion. In fact, also Sener et al. (2012), Bali and Theodossiou (2007), Abad and Benito (2013), Chen et al. (2011), Mittink and Paolella (2000) support this thesis.

The results obtained by comparing simple Monte Carlo simulation with other methods show that other methods are generally able to provide better results. The papers of Abad and Benito (2013), Huang (2009), Tolikas et al. (2007), Bao et al. (2006) testify this conclusion. Bao et al. (2006) anzlyze a series of Asian stocks belonging to five different markets that suffered from the 1997–1998 financial crisis. Simple Monte Carlo simulation performs equally or worse than the other methods taken into account.

When compared to a variety of alternative parametric models and basic historical simulation-based techniques, Chan and Gray (2006) found that the EVT delivered the adequate unconditional and conditional coverage. The authors compare the performance of different models in delivering an adequate VaR estimate in the electricity market. It is important to underline the fact that, even more than with traditional financial assets, simple methods are likely to produce poor Value at Risk estimates in such a market in which returns are highly volatile and show seasonality in the mean as well as in the variance, leverage effects and volatility clustering, and abnormal levels of skewness and kurtosis. The data analyzed are daily aggregated electricity spot prices from five international power markets: Victoria (Australia), NordPool (Scandinavia), Alberta (Canada), Hayward (New Zealand) and PJM (US). The models examined in the paper are the following: simple historical simulation approach, different extensions of auto-regressive AR (7) models, one assuming constant variance; one in which variance is obtained through bootstrapping, one in which the error term is assumed to follow an EGARCH process, both under normally and t-distributed error terms; the so-called AR-EGARCH-EVT, in which Extreme Value Theory is applied to the GARCH residuals. The full data sample is divided into an in-the-sample period and an out-of-the-sample period over which VaR performance is evaluated. The more sophisticated AR-EGARCH-EVT approach is found to dominate in markets where the distribution of returns is characterized by high levels of skewness, kurtosis, and volatility, while simple historical simulation is acceptable in markets where the pattern of returns used to build the model was not different to the one in the evaluation period.

Again on Extreme Value Theory, Kuester et al. (2006) extended the EVT framework set up by McNeil and Frey (2000) and argued that a hybrid method, combining a heavy-tailed (such as t-distributed) GARCH filter with an extreme value theory based approach, performed better than the alternatives such as historical simulation and fully parametric approaches.

Swami, Pandey and Pancholy (2016) conduct a study on Value at Risk estimation with foreign exchange rates in India. In particular, they take data from 1999 to 2013 for three exchange rates: Indian Rupee (INR) / US Dollar (USD), INR / Pound and INR / EUR for a theoretical portfolio value equal to 200 million Rupees, with weights, respectively, equal to 50 percent, 25 percent, and 25 percent in each of the previously mentioned foreign currencies. Then, they employ different approaches to estimate VaR (ranging from the variance-covariance method to non-parametric approaches) and adopt backtesting techniques to evaluate each of them; backtesting is carried out for each year from 2000 to 2013 by using the preceding 1-year data. Histograms and Q-Q plots show that even FX rates are not normally distributed. As a matter of fact, the results of the paper reveal that assuming normality will not provide a robust risk assessment; when assuming t-distributed rates, the assessment is instead more appropriate.

Value at Risk focuses on the risks of a single bank in isolation, without taking into account systemic effects. As I already mentioned in the previous chapter, though, the Global Financial Crisis showed that the difficulties of some institutions can create spillovers that hurt other banks that would otherwise be safe. In order to take this into account, Adrian and Brunnermeier (2008) proposed a measure called Conditional VaR ΔCoVaR (it is important to note that even though also the Expected Shortfall measure can be called Conditional VaR, CVaR, these are two different concepts). The purpose is to study how the VaR of a particular bank is affected by negative externalities created by other banks in the financial system. It is defined as the change in an institution's VaR conditional on a bank being in a situation of crisis or distressed evaluated with respect to its median state. In other words, ΔCoVaR is the difference between CoVaR conditioned on the institution's state of distress and the CoVaR conditioned on its median state. It is thus a systemic risk indicator that comoves with the distress of a certain bank.

Giraldi and Ergun (2013) change the original definition of financial distress of a financial institution from being exactly at its VaR to having returns equal or below the VaR, which allows them to study even more severe distress situations. The authors define the standard case as all the events for which returns are one-standard deviation about the mean event. The analysis takes into account data in the time period from 2000 to 2008 for a wide number of institutions (74) belonging to four industry groups: depositories, insurers, broker-dealers, and others (such as government sponsored enterprises), while the benchmark adopted as proxy for systemic returns is the US Dow Jones Financials Index (DJUSFN). The novelty of the research is that the authors also investigate the link between institutions' characteristics such as size, leverage and levered beta and their respective contributions to systemic risk by running panel regressions. The result the authors obtain is that the factor, among the three, that most correlates with CoVaR is the equity beta. The pre-crisis analysis shows that the systemic risk of all industry areas increased substantially during the 12 months before June 2007. CoVaR estimates based on the assumption of normally distributed returns fail the unconditional coverage test, while t-distributed returns pass both the unconditional and the conditional coverage test. Another interesting result of the study is the low correlation between a single institution's VaR and CoVaR in the cross-section analysis; thus, from a regulatory point of view, the risk that each institution actually poses to the financial system may not be taken into consideration if capital requirements are only based on an institution's VaR rather than also taking into account the CoVaR of each institution.

Nolde, C. Zhou and M. Zhou (2022) performed a study by adapting extreme value theory to semi-parametric estimation of CoVaR. In their research, the authors consider the returns of 14 financial institutions while the S&P 500 index is used as a proxy of the systemic returns. The time period spans from 2000 to 2021, consisting of 5535 daily closing price records for each time series. Although CoVaR can theoretically be backtested in the same way as VaR, conditioning on an institution's losses exceeding the VaR estimate creates a practical difficulty to obtain statistically significant results when performing backtesting due to the substantial reduction in the size of the testing data set. However, the proposed CoVaR estimator is shown to be consistent. The authors also propose a way to model tail dependence in a more flexible and less computationally intensive way than alternative models. A limitation of their approach, though, is that it is based on the assumption of tail-dependence, which, while reasonable for financial time series, limits the overall applicability of the model.

4 Empirical Study

In this section, I describe the methodology that I have employed to carry out the empirical study and show the relevant results that I have obtained. Specifically, I divide the chapter into four subsections: in the first I present the data that I have used to build the portfolio, in the second I report the sample period taken into account and the framework for the evaluation of the models, in the third I detail the construction of the above mentioned models, and in the fourth I show the results of the study.

4.1 Data

To carry out this study, I built a portfolio composed of five different asset classes, specifically the US stock index NASDAQ Composite (ˆIXIC) that includes almost all the stocks listed on the Nasdaq stock exchange and is heavily weighted towards companies in the IT sector, a fixed income security issued by the government of Italy on 03/02/1997 and with maturity 01/11/2026, paying a fixed semiannual coupon equal to 7.5%, the commodity index PIMCO CommoditiesPLUS Strategy Fund (PCLIX), plus the Bitcoin/USD exchange rate and the EUR/USD exchange rate[13]. I chose this composition in order to have a well-diversified (not only in terms of asset class but also in terms of geopgraphy) and liquid portfolio, with a US stock index and a European fixed income security, plus commodities, FX and a cryptocurrency.

I obtained dollar daily adjusted closing prices (not taking into account splits and dividends) for the five assets for the period from 17/08/2020 to 18/08/2022, thus obtaining 500 observations. It is important to mention that this time period has to be considered as a crisis period, because of the Covid pandemic and the Ukrainian war, that highly increased the volatility of returns, and, for some assets, also deteriorated their financial

[13] I gathered all the financial data through www.finance.yahoo.com, except the data on the Italian government bond that was found on the Refinitiv platform. I have taken dollar prices for all the securities. The price of the Italian bond was expressed by default as a pure number as ratio of the price over the face value of the bond (100 euros); however, I downloaded the dollar price of the security thanks to the conversion that Refinitiv allows.

performance. I constructed a $10 million portfolio by considering the fact that with riskier securities (such as the Bitcoin) having a greater proportion in the portfolio would be too risky, so the portfolio was built in the following way: 25% of the total amount invested in ˆIXIC, another 25% invested in the Italian bond, then 20% in the PCLIX, 10% in the Bitcoin/USD and 20% in the EUR/USD exchange rate.

4.2 Evaluation Framework

The 2-year data that I obtained were split so that the first 250 observations served as in-the-sample to generate VaR estimates that were then backtested on the second year of data. For all asset classes, the daily VaR was constructed through a rolling window that included the last 250 observations. In other words, I do not have the Value at Risk for the first 250 observations, which are needed for the estimation of the VaR for the 251th business day and on. VaR was tested at 1% confidence level one-step ahead, so that for a robust VaR model I expected 2 or 3 exceptions over the 250 out-of-the-sample observations. I performed both the Unconditional Coverage Test (Kupiec test) and the Conditional Coverage Test (Christoffersen test) first for each security and then at portfolio aggregate (of course I did not perform the Conditional Coverage test when there were no exceptions at asset level or portfolio level). At aggregate portfolio level, I performed backtesting for two different measures: the portfolio Net VaR, which takes into account the benefit of diversification, and the portfolio Gross VaR, simply obtained by summing the individual VaRs of the five securities. The difference between Gross VaR and Net VaR allows to quantify in monetary terms the diversification effect.

4.3 VaR models

For the purpose of the study, I built VaR models through simulation methods, both historical simulation and Monte Carlo simulation to evaluate the performance of these approaches in the assessment of market risk for a portfolio of securities belonging to five different asset classes. After having computed the Value at Risk estimates, I performed the backtesting to assess the validity of the models. In the following paragraphs I will detail the procedure I followed to develop the VaR simulations.

4.3.1 Historical Simulation

In this subsection I am going to describe the methodology I used to build the Value at Risk employing the historical simulation approach. First I imported the relevant daily prices, I then computed the log-returns of each security by taking the natural logarithm of the price at time t divided by the price at time $t-1$ and the Profit and Loss (PnL) at time t by multiplying the value of the portfolio at time t and the log-return observed a time t. I estimated the 1% daily VaR for all the securities by taking the first percentile of the PnLs with the Excel formula =PERCENTILE.EXC, with a rolling window of 250 PnLs. I adopted equally weighted historical simulation, thus assigning the same proportion of relevance to more and less recent observations. Once I had obtained the VaR level for the 250 out-of-the-sample observations, I checked whether at each t an exception occurred or not by comparing the daily VaR with the daily Profit and Loss; I used an indicator function to count as 1 when the exception occurred (loss greater than the VaR), and 0 when the exception did not occur. With this information, I performed the Unconditional Coverage Test (Kupiec test) to count the number of exceptions and have an understanding of whether historical simulation was able to adequately assess market risk (i.e. whether the percentage of such exceptions was in line with the desired percent α. Moreover, I carried out the Conditional Coverage Test (Christoffersen test) to check whether the pattern of exceptions is serially dependent (i.e. whether exceptions are clustered or not). These tests were then performed at the aggregate portfolio level to obtain two measures: the portfolio Net VaR and the portfolio Gross VaR. To obtain the Net VaR estimate, I applied the percentile function to the daily portfolio PnL (obtained by summing the PnLs of the individual securities); the Gross VaR estimate was instead obtained by summing the VaR figures of the five instruments. Afterwards, I performed the same procedure of backtesting that I used for the individual securities.

4.3.2 Monte Carlo Simulation

In this part I show the steps that I have followed in order to build VaR estimates under the Monte Carlo simulation for the asset classes that compose the portfolio. Besides the general concept of Monte Carlo simulation, I have assumed that different processes regulate the evolution of the path of the prices of different assets.

ˆIXIC Stock Index

For the stock index NASDAQ Composite (ˆIXIC), I have first of all computed daily log-returns and daily PnL in the same way I described in the previous section. Then, I assumed that the evolution of the log-return of the stock index follows a Geometric Brownian Motion (GBM), such that:

$$S_{t+1} = S_t \exp\left(\mu - \frac{\sigma}{2}\right)^2 \Delta t + \sqrt{\Delta t}\epsilon_{t+1} \tag{30}$$

where S_{t+1} is the simulated price of the index one-day ahead, S_t is the price today, μ and σ are the average log-return and standard deviation computed on a rolling window of 250 days, Δt is the time interval, namely 1/250, and ϵ_{t+1} is the random shock distributed as N(0,1).

Therefore, I first computed μ, σ and $(\mu - \frac{\sigma}{2})^2$; then I simulated 2100 iterations of a random standard normal distribution with the Excel command =NORMSINV(RAND()). Then I simulated the log-return arising from the previously simulated realization of the standard normal; and the implied portfolio value assuming the simulated realization. The last step to compute the Value at Risk estimate was to obtain two measures for the 2100 simulations of the portfolio value in $t + 1$: the average portfolio value and the portfolio value at the 1st percentile (needed for the VaR at 1%). The relevant VaR output is the difference between the two measures. Once I obtained the VaR estimate, I backtested the result with the Unconditional Coverage Test and Conditional Coverage Test in the same way as I described in the previous paragraph.

Italian Government Bond ISIN T0001086567

With respect to the fixed income security, I applied the Vasicek model in order to simulate the evolution of interest rates, and from that, the evolution of the bond price so as to obtain the VaR. I will now detail the steps I followed to obtain the VaR estimate. First, I compute the log-return and PnL as usual. Then, I compute the yield (y_t) of the bond for each t with the Excel function =YIELD. Once I had y_t, I found the corresponding interest rate r_t with the Goal Seek functionality. In order to estimate the path evolution of interest rates I assumed a Vasicek model to simulate them of the type:

$$r_{t+1} = r_t + \alpha(\beta - r_t)\Delta t + \sigma\sqrt{\Delta t}\epsilon_{t+1} \tag{31}$$

where α and β are, respectively, the speed of reversion and the long-term mean of interest rates towards which they tend to revert and ϵ_{t+1} is a random shock distributed as a standard normal. $\alpha(\beta - r_t)$ is a drift factor that represents the expected instantaneous change in the interest rate at time t. I found the coefficients of α and β through the regression of the following AR (1) process:

$$r_t = \alpha + \beta r_{t-1} + \epsilon_t \tag{32}$$

with standardized residuals $\epsilon_t \sim$ N(0,1).

Below, I show the summary of the output of the regression performed to find the coefficients of α and β under equation (32).

After having found the coefficients α (the intercept) and β (the coefficient of the X variable), I carried out 2100 simulations of the Vasicek model (31) so as to simulate the possible evolutions of interest rates during the second year of data. Once I had the iterations of the interest rate paths, I retrieved the average and 1st percentile yield and

SUMMARY OUTPUT

Regression Statistics	
Multiple R	0.98950915
R Square	0.97912835
Adjusted R Square	0.97904385
Standard Error	0.00011187
Observations	249

ANOVA

	df	SS	MS	F	Significance F
Regression	1	0.000145006	0.000145	11587.24	0.00
Residual	247	0.00	0.00		
Total	248	0.000148097			

	Coefficients	Standard Error	t Stat	P-value
Intercept	0.00040697	0.000431045	0.94415	0.346016
X Variable 1	0.99140337	0.009210018	107.644	0.00

Figure 1: Regression Summary

then the average and 1st percentile price of the bond. Given the price levels I had found, I retrieved the average and 1st percentile portfolio value for the dates on which then I performed the backtesting. Taking the difference between the average and 1st percentile portfolio value I found the Value at Risk measure; finally, I performed the Kupiec test and the Christoffersen test in the same way I described previously.

It must be pointed out that, for this type of instrument there was an additional source of risk which other securities did not have. In fact, given that the other asset classes were quoted in dollars, I have transformed the price of the Italian Government bond, quoted in euros, into dollar prices, so as to have a meaningful portfolio VaR outcome.

Commodities Index PCLIX

In order to simulate the price evolution of a commodity index I assumed that, first of all, it has a component of mean reversion towards a long-term mean (I have used the last 4 year average price and assumed a mean reversion factor of 0.6%). Specifically, I assumed that the price of the commodity index at time $t+1$ followed the following process:

$$P_{t+1} = P_t \exp\left(\sigma\epsilon_{t+1}\right) + k(\mu - P_t) \tag{33}$$

where k is the mean reversion factor and μ is the long-term mean. I then simulated 2100 iterations of P_{t+1} through which I computed the average portfolio value and the 1st percentile portfolio value and followed the steps already presented to compute the VaR for the commodity index.

BTC/USD

I decided to include in the portfolio also the price of a cryptocurrency such as Bitcoin (BTC) computed against the price of US Dollar (USD). I assumed the evolution of the Bitcoin price to be governed by a Geometic Brownian Motion process, similarly to the stock index included in the portfolio. Therefore, the procedure I implemented to compute the VaR and then backtest its performance was the same I described for ˆIXIC index.

EUR/USD

For the last instrument included in the portfolio, namely the EUR/USD exchange rate, I assumed a similar process to the one for the BTC/USD. However, I assumed that the shock is not normally distributed but it is distributed as a t-student with 1 degree of freedom (thus incorporating fatter tails in the unpredictable shock). For the other steps, I applied what I did for the BTC/USD exchange.

4.4 Results

4.4.1 Historical Simulation

Below, I show the graphical representation of the daily Profit and Loss and the daily VaR estimates under the historical simulation approach of the five assets included in the portfolio plus the portfolio PnL and Net and Gross VaR. Then, I present the outcome of the backtesting, and finally, I discuss the results obtained using the historical simulation approach.

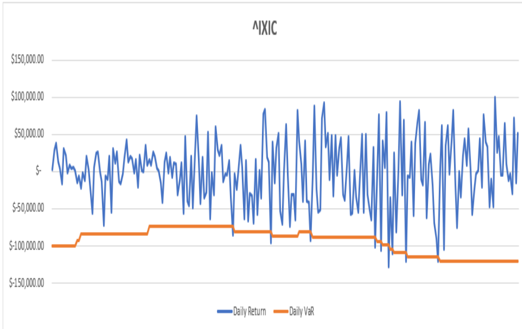

Figure 2: Daily VaR and Daily PnL for the ˆIXIC Index

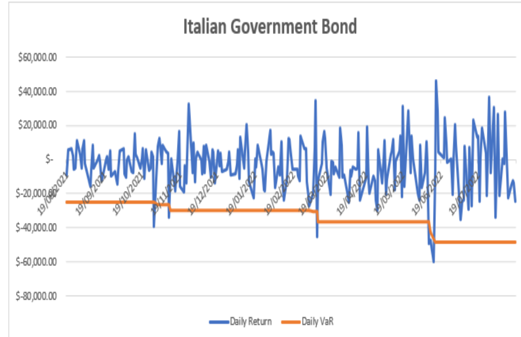

Figure 3: Daily VaR and Daily PnL for the Italian Government bond

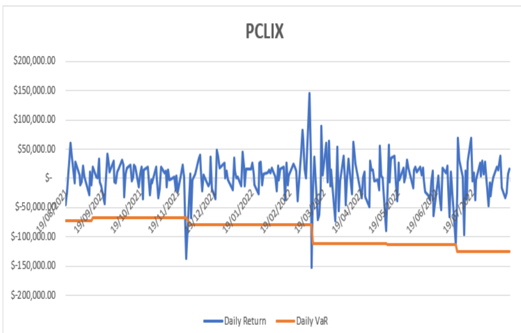

Figure 4: Daily VaR and Daily PnL for the PCLIX Commodities Index

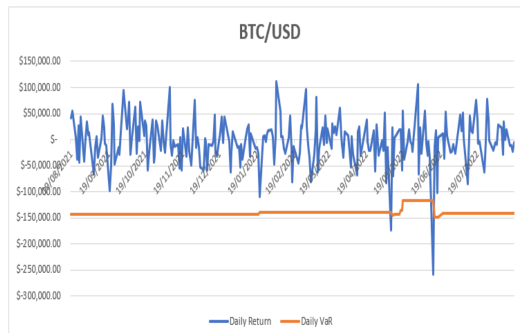

Figure 5: Daily VaR and Daily PnL for the Bitcoin price against the USD

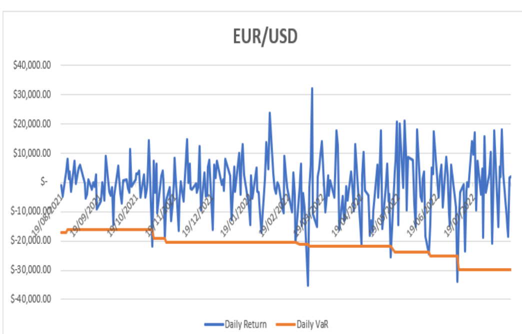

Figure 6: Daily VaR and Daily PnL for the EUR/USD exchange rate

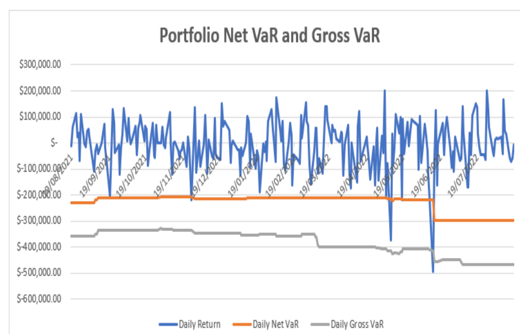

Figure 7: Daily VaR and Daily PnL for the whole portfolio

The tests on the validity of the VaR estimate under the historical simulation approach show that for some securities the outcome is not satisfactory and it does not meet the criteria to pass the Unconditional and the Conditional Coverage Tests. In fact, the amount of

n1	9
n	250
n0	241
Kupiec test	**10.2290**
p value	**0.00138**

n00	227
n01	9
n10	9
n11	0
pi0	0.03814
pi1	0
pi	0.03673
Christoffesen test	**114.36350**
p value	**0.00000**

Figure 8: Backtesting results for the ˆIXIC Index

n1	6
n	250
n0	244.0135594
Kupiec test	**3.53126**
p value	**0.06022**

n00	239
n01	4
n10	4
n11	2
pi0	0.01646
pi1	0.33333
pi	0.02410
Christoffesen test	**80.77562**
p value	**0.00000**

Figure 9: Backtesting results for the Italian Government bond

n1	3
n	250
n0	247
Kupiec test	**0.09494**
p value	**0.75799**

n00	243
n01	3
n10	3
n11	0
pi0	0.01220
pi1	0
pi	0.01205
Christoffesen test	**47.67854**
p value	**0.00000**

Figure 10: Backtesting results for the PCLIX Commodities Index

n1	2
n	250
n0	248
Kupiec test	**0.10844**
p value	**0.74193**

n00	245
n01	2
n10	2
n11	0
pi0	0.008097166
pi1	0
pi	0.008032129
Christoffesen test	**33.90548**
p value	**0.00000**

Figure 11: Backtesting results for the Bitcoin price against the USD

n1	6
n	250
n0	244
Kupiec test	**3.55535**
p value	**0.05935**

n00	237
n01	6
n10	6
n11	0
pi0	0.02469
pi1	0.00000
pi	0.02410
Christoffesen test	**83.69569**
p value	**0.00000**

Figure 12: Backtesting results for the EUR/USD exchange rate

n1	4
n	250
n0	246
Kupiec test	0.76914
p value	0.38048

n00	241
n01	4
n10	4
n11	0
pi0	0.01633
pi1	0.00000
pi	0.01606
Christoffesen test	**60.41169**
p value	**0.00000**

Figure 13: Backtesting results for the Net VaR of the portfolio

exceptions (represented by n1 in the tables above) for the NASDAQ Composite Index and for the Italian government bond are too high compared to the α of the test: in fact, the institution having such instruments would fall in the Yellow Zone under the Basel Accord penalty zones. For other securities, however, the VaR estimates performs well, in particular for the PCLIX Commodities Index and for the Bitcoin, for which 3 and 2

n1	**1**
n	250
n0	249
Kupiec test	1.17649
p value	0.27807

n00	247
n01	1
n10	1
n11	0
pi0	0.00403
pi1	0.00000
pi	0.00402
Christoffesen test	**18.54229**
p value	**0.00002**

Figure 14: Backtesting results for the Gross VaR of the portfolio

exceptions, respectively, are in line with what is expected in terms of α. Also for the aggregate portfolio the 4 exceptions in a year would locate them in the Green Zone of the Basel Accords, which, in a turbulent time period such as the one under analysis is a satisfactory performance for the VaR estimate. This is crucial because with a good portfolio diversification, a bank is able to correctly assess its market risk even in a period of stress and even if for some securities inside the portfolio the risk estimation is not optimal.

It is important to note, however, the amount of the Value at Risk size for the different assets in the portfolio. For instance, even though the VaR for the Bitcoin only has two exceptions in a year, it has a magnitude in terms of capital requirement much greater than the other securities. As a matter of fact, the average VaR level is around 14% of the initial investment, while for the other securities it ranges between 1% and 4%. This means that an institution with Bitcoin in its portfolio would fall in the green zone, thus would not be subject to penalties of additional capital requirements, but the level of capital itself would be much greater than that for other securities. This huge amount of capital is justified by the large swings in the price that the cryptocurrency realized in the period under analysis, with losses up to 257.000 dollars in one day with an investment of 1 million dollars.

I have presented the fact that in terms of Unconditional Coverage Tests, the VaR performed well for some asset classes, quite poorly for others, even though, considering the stress period, the scenario was not catastrophic. As for the Conditional Coverage Test, however, none of the assets analyzed, nor the portfolio aggregate was able to pass the Christoffersen test. With exceptions clustered, as the results of the test show, the likelihood to see the capital wiped out increases. Therefore, in terms of the dependence of exceptions, the performance is absolutely not adequate.

4.4.2 Monte Carlo Simulation

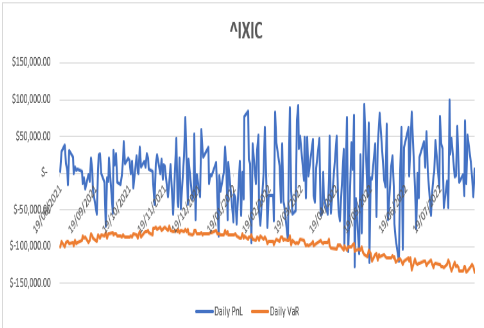

Figure 15: Daily VaR and Daily PnL for the ˆIXIC Index

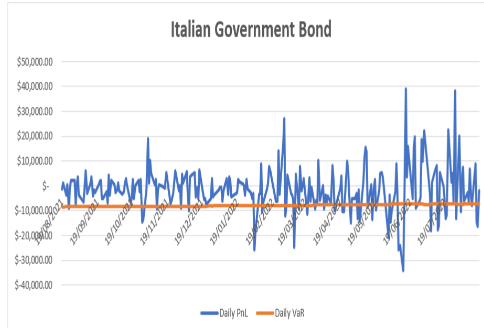

Figure 16: Daily VaR and Daily PnL for the Italian Government bond

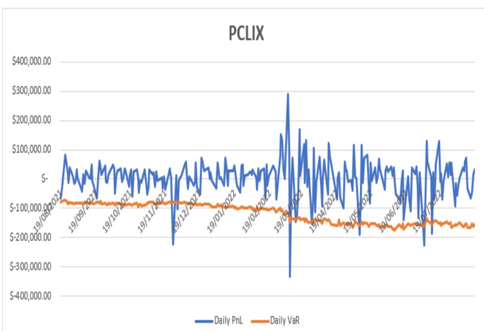

Figure 17: Daily VaR and Daily PnL for the PCLIX Commodities Index

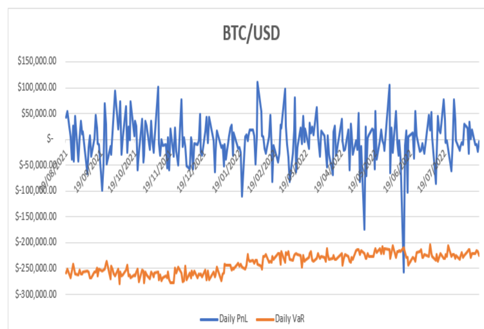

Figure 18: Daily VaR and Daily PnL for the Bitcoin price against the USD

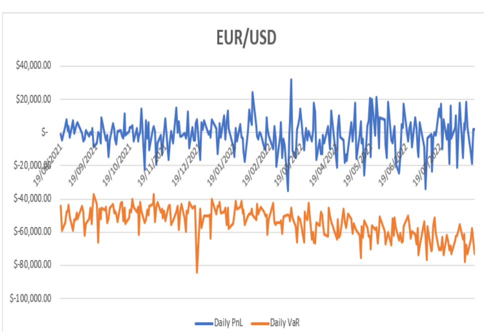

Figure 19: Daily VaR and Daily PnL for the EUR/USD exchange rate

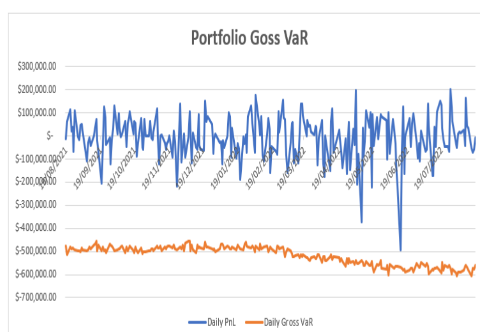

Figure 20: Daily VaR and Daily PnL for the whole portfolio

n1	7
n	250
n0	243
Kupiec test	**5.49699**
p value	**0.01905**

n00	235
n01	7
n10	7
n11	0
pi0	0.02893
pi1	0.00000
pi	0.02811
Christoffesen test	**94.50190**
p value	**0.00000**

Figure 21: Backtesting results for the ˆIXIC Index

n1	**42**
n	250
n0	208
Kupiec test	**164.66487**
p value	**0**

n00	172
n01	31
n10	31
n11	10
pi0	0.15271
pi1	0.24390
pi	0.16803
Christoffesen test	**330.47872**
p value	**0.00000**

Figure 22: Backtesting results for the Italian Government bond

n1	8
n	250
n0	242
Kupiec test	**7.73355**
p value	**0.00542**

n00	232
n01	8
n10	8
n11	0
pi0	0.03333
pi1	0.00000
pi	0.03226
Christoffesen test	**104.75765**
p value	**0.00000**

Figure 23: Backtesting results for the PCLIX Commodities Index

n1	**1**
n	250
n0	249
Kupiec test	**1.17649**
p value	**0.27807**

n00	246
n01	1
n10	1
n11	0
pi0	0.00405
pi1	0.00000
pi	0.00403
Christoffesen test	**18.53018**
p value	**0.00002**

Figure 24: Backtesting results for the Bitcoin price against the USD

n1	0

Figure 25: Backtesting results for the EUR/USD exchange rate

n1	0

Figure 26: Backtesting results for the Gross VaR of the portfolio

In this subsection I have followed the schema adopted in the previous one: I first showed the graphical representation of the daily Profit and Loss and the daily VaR estimates under the Monte Carlo simulation approach of the five assets included in the portfolio plus the portfolio PnL and Net and Gross VaR. Then, I showed the outcome of the backtesting carried out for the five assets and for the whole portfolio. I will now discuss the results obtained using the Monte Carlo simulation approach.

The results obtained adopting this second type of simulation approach are controversial as to the validity of this approach. In fact, for some assets, Monte Carlo simulation performs well also in the period of stress analyzed. In fact, there are no exceptions for the EUR/USD exchange rate and only one for the Bitcoin. However, it must be noted that capital requirements for the EUR/USD exchange rate and for the Bitcoin price are much greater than those that I ended up under the historical simulation approach; this may signal that the level of capital put aside is too high. On the other hand, the performance for the fixed income instrument, namely the Italian bond, shows a catastrophic performance. In fact, out of 250 observations, there are 42 exceptions, which would, of course, locate the bank into the red zone under the Basel Accords. It is clear that the risk forecast is not trustworthy and capital requirements based on this estimate would be wiped out easily. Note that during the period under analysis the price of the bond constantly declined but volatility was not high, therefore the simulation generates values that are not high enough to cover from the risk of this price drop.

Similarly to the results that I obtained under the historical simulation approach, exceptions are dependent and the criteria to pass the Conditional Coverage test are not passed, which again poses a problem for the solvency of a bank that experiences large, concentrated, losses.

It is critical to note that Value at Risk numbers materially differ between the two simulation approaches: in particular, the VaR estimate for the Bitcoin was an average of around $139.000, while if obtained through the Monte Carlo simulation was on average $239.000; similarly, also the Gross portfolio VaR obtained through Monte Carlo simulation ended up to be significantly greater than that obtained via historical simulation. These findings indicate that it may be highly misleading to compare VaR estimates across financial institutions that report numbers based on different methods.

5 Conclusions

The main aim of this work is to study and compare the validity of simulation approaches (historical and Monte Carlo) to VaR estimation in a period of stress for financial markets (from August 2020 to August 2022, when the world was hit by the Covid pandemic and by the war in Ukraine). I built a portfolio composed of five different asset classes, namely a US stock index, an Italian Government bond, a commodities index, the FX EUR/USD exchange rate, and the Bitcoin cryptocurrency and evaluated the accuracy of market risk estimation by adopting two backtesting procedures: the Unconditional Coverage test, used to study whether the intended confidence level at which VaR is computed is actually captured by the model, and the Conditional Coverage test, that takes into account the time distribution of exceptions.

I find that, if taking into account the securities that compose the portfolio one by one, neither of the two methods are able to prove their validity for all the instruments in such a period of stress. In particular, exceptions are too numerous for the ˆIXIC Index and for the Italian Government bond under both approaches, and for the latter security the outcome is catastrophic with the Monte Carlo simulation. However, when I analyzed the VaR estimate at portfolio aggregate, I witnessed to an improvement in the risk assessment: in fact, the Net VaR computed with historical simulation gave rise to 4 exceptions, which is an acceptable result, in particular in a stress period, and would locate the bank in the green zone under the Basel Accord penalty zones. This result is of important because it grants that if banks are able to diversify their portfolios, they are able to generate appropriate VaR estimates even in a period of stress and even if for some securities inside the portfolio the risk estimation is not optimal. The Gross VaR estimated under historical simulation and Monte Carlo simulation generated, respectively, 1 and 0 exceptions, but in this case one might argue that the level of capital put aside is too high and fails to consider the benefits of diversification into different asset classes. In general, historical simulation performed better than Monte Carlo simulation as for the Unconditional Coverage test. In terms of conditionality of exceptions, neither the individual securities nor the portfolio aggregate are able to pass the Conditional Coverage test and this is dangerous because clustered exceptions may lead to the erosion of the capital of a bank.

6 References

Abad, P. & Benito, S. (2013), A Detailed Comparison of Value at Risk Estimates, *Mathematics and Computers in Simulation*, volume 94, 258-276

Adrian, T. & Brunnermeier, M. (2013), A Detailed Comparison of Value at Risk Estimates, *The American Economic Review*, volume 106, issue 7, 1705-1741

Angelidis, T. & Benos, A. & Stavros, A. (2007), A robust VaR model under different time periods and weighting schemes, *Review of Quantitative Finance and Accounting*, volume 28, issue 2, 187-201

Aramonte, S. & Rodriguez, M. & Wu, J. (2013), Dynamic factor Value-at-Risk for large heteroskedastic portfolios, *Journal of Banking & Finance*, volume 37, issue 11, 4299-4309

Ashley, R. & Randal, V. (2009), Frequency Dependence in Regression Model Coefficients: An Alternative Approach for Modeling Nonlinear Dynamic Relationships in Time Series, *Econometric Reviews*, volume 28, issue 1-3, 4-20

Bali, R. & Theodiossou, V. (2007), Risk Measurement Performance of Alternative Distribution Functions, *Journal of Risk & Insurance*, volume 75, issue 2, 411-437

Bao, Y. & Lee, H. & Saltoglu, B. (2006), Evaluating predictive performance of value at risk models in emerging markets: a reality check, *Journal of Forecasting*, issue 25, 101–128

Barone-Adesi, G. & Giannopoulos, K. & Vosper, L. (1999), VaR without correlations for portfolios of derivative securities, *The Journal of Futures Markets*, volume 19, issue 5, 583-602

Beder, T. (1995), VAR: Seductive but Dangerous, *Financial Analysts Journal*, volume 51, issue 5

Berkovitz, J. & O'Brien, J. (2002), How Accurate Are Value-at-Risk Models at Commercial Banks?, *The Journal of Finance*, volume 57, issue 3, 1093-1111

Billio, M. & Pellizon, L. (2000), Value-at-Risk: a multivariate switching regime approach, *Journal of Empirical Finance*, volume 7, issue 5, 531-554

Black, F. (1976), Studies of Stock Price Volatility Changes, *Proceedings of the 1976 Meeting of the Business and Economic Statistics Section, American Statistical Association*

Chan, K. F. & Gray, P. (2006), Using Extreme Value Theory to Measure Value-at-Risk for Daily Electricity Spot Prices, *International Journal of Forecasting*, volume 22, 283-300

Engle, R. (2002), Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models, *Journal of Business & Economic Statistics*, volume 20, issue 3, 339-350

Engle, R. & Kroner, K. (1995), Multivariate Simultaneous Generalized Arch, *Econometric Theory*, volume 11, issue 1, 122-150

Giraldi, G. & Ergun, A. (2013), systemic risk measurement: Multivariate GARCH estimation of CoVaR, *Journal of Banking & Finance*, volume 37, issue 8, 3169-3180

Guermat, C. & Harris, R. (2002), Forecasting value at risk allowing for time variation in the variance and kurtosis of portfolio returns, *International Journal of Forecasting*, volume 18, issue 3, 409-419

Gonzalez-Rivera, G. & Lee, T. H. & Mishra, S. (2004), Forecasting volatility: A reality check based on option pricing, utility function, value-at-risk, and predictive likelihood, *International Journal of Forecasting*, volume 20, issue 4, 629-645

Hansen, B. (1994), Autoregressive Conditional Density Estimation, *International Economic Review*, volume 35, issue 3, 705-730

Harvey, C. R. & Siddique, A. (1999), Autoregressive Conditional Skewness, *Journal of Financial and Quantitative Analysis*, volume 34, 465-487

Hendricks, D. (1996), Evaluation of Value-at-Risk Models Using Historical Data, *FRBNY Economics Policy Review*, volume 2, issue 1, 39-70

Jondeau, E. & Rockinger, M. (2001), Conditional Dependency of Financial Series: An Application of Copulas, *Groupe HEC-Department Finance Economics and Banquede France-Economic Study and Research Division*

Kuester, K. & Mittnik, S. & Paolella, M. (2006), Value-at-Risk Prediction: A Comparison of Alternative Strategies, *Journal of Financial Econometrics*, volume 4, issue 1, 53-89

Li, H. & Fan, X. & Li, Y. & Zhou, Y. & Jin, Z. & Liu, Z. (2011), Approaches to VaR, *Stanford University*

Marimoutou, V. & Raggad, B. & Trabelsi, A. (2009), Extreme Value Theory and Value at Risk : Application to Oil Market, *Energy Economics*, volume 31, issue 4, 519-530

Niguez, T. (2008), Volatility and VaR Forecasting in the Madrid Stock Exchange, *Spanish Economic Review*, volume 10, issue 3, 169-196

Nozari, M. & Raei, S. M. & Jahanguin, P. & Bahramgiri, M (2010), A comparison of heavy-tailed estimates and filtered historical simulation: evidence from emerging markets, *International Review of Business Papers* volume 6, issue 4, 347–359

Pagan, A. & Schwert, G. W. (1990), Alternative models for conditional stock volatility, *Journal of Banking & Finance*, volume 45, issues 1-2, 267-290

Pérignon, C. & Smith, R. (2010), The level and quality of Value-at-Risk disclosure by commercial banks, *Journal of Banking & Finance*, volume 34, issue 2, 362-377

Polanski, A. & Stoja, E. (2009), Incorporating higher moments into value-at-risk forecasting, *Journal of Forecasting*, volume 29, issue 6, 523-535

Pritsker, M. (1997), Evaluating Value at Risk Methodologies: Accuracy versus Computational Time, *Journal of Financial Services Research*, volume 12, 201-242

Pritsker, M. (2006), The hidden dangers of historical simulation, *Journal of Banking & Finance*, volume 30, issue 12, 561-582

Raaji, G. D. & Raunig, B. (1998), A comparison of Value At Risk Approaches and their implications, *Focus on Austria*, volume 4, 57-71

Tolikas, I. & Gettinby (2009), Modelling the distribution of the extreme share returns in Singapore, *Journal of Empirical Finance*, volume 16, issue 2, 254-263

Swami, O. & Pandey, S. & Pancholy, P. (2009), Value-at-Risk Estimation of Foreign Exchange Rate Risk in India, *Asia-Pacific Journal of Management Research and Innovation*, volume 12, issue 1

Trenca, I. (2009), The use in banks of VaR method in risk management, *Scientific Annals of the 'Alexandra Ioan Cuza'*, volume 56, 186-196

Zikovic, S. & Aktan, B. (2009),Global financial crisis and VaR performance in emerging markets: A case of EU candidate states-Turkey and Croatia, *Proceedings of Rijeka School of Economics*, volume 27, issue 1, 149-171

Resti, A. & Sironi, A. (2007), *Risk Management and Shareholders' Value in Banking: From Risk Measurement Models to Capital Allocation Policies*, Wiley, 1st Edition

Summary

The main objective of this work is to present and discuss the results of the analysis that I carried out to compare different approaches to Value at Risk (VaR), one of the main tools that banks adopt for risk management purposes.

With the objective of preserving financial stability, regulators in most areas of the world have decided to impose that banks at all times have layers of capital that work as a buffer in case banks face unexpected losses, and VaR is the measure that most financial institution have been adopting in the last decades to estimate the level of capital they must put aside.

The structure of the thesis is organized as follows: three chapters that follow this brief introduction and that constitute the main body of the dissertation, and a final section containing the conclusions.

The first chapter contains a theoretical analysis of what is Value at Risk. Here I begin with a general overview of banks risk management practices with a focus on market risk. Then, I proceed with a deep dive on the VaR concept, its merits and its drawbacks and the regulatory foundations on which VaR is rooted. I then present the three standard approaches adopted to estimate VaR, namely the Variance-Covariance approach, the historical simulation and the Monte Carlo simulation approach. I also discuss different models for volatility estimation, crucial to build appropriate VaR models, ranging from simple moving average methods to more sophisticated GARCH models. Finally, I discuss what is backtesting, why it is adopted to review VaR performance, and what type of tests can be performed.

The second chapter includes a dissertation of the existing literature on Value at Risk relevant for the study that follows. After the presentation of a work that investigates what approaches financial institutions actually adopt, I discuss the results obtained by researchers on the VaR topic. The most investigated aspects are what models perform best for volatility estimation and for VaR estimation.

In the third chapter, I present the empirical study that I have performed in order to assess the performance of the historical simulation approach and the Monte Carlo simulation approach for a portfolio composed of five different asset classes in a stress period for financial markets (from August 2020 to August 2022). I begin by describing the construction of the portfolio, from the data of the securities that compose the portfolio to the weights attached to each instrument; then, I discuss the evaluation framework adopted, in particular how I have divided the data into a rolling window for in-the sample model construction and out-of-the sample for model estimation, and, for the latter, the backtesting that I have performed. I have then proceeded with a detailed analysis of how I estimated the Value at Risk levels under the two simulation approaches. With respect to historical simulation, I cut the relevant daily PnL distribution at the desired percentile (1%), while for the Monte Carlo simulation I have simulated price evolutions assuming different processes for the different asset classes. I then computed the Gross VaR of the portfolio under the two approaches, and the Net VaR of the portfolio under the historical simulation approach. To conclude, I showed the results of the analysis, focusing on how well the VaR models performed.

Value at Risk defines what is the maximum loss of a position that will be sustained given a certain confidence level within a defined horizon of time. It has the merit to provide a single number which represents the exposure to market risk in monetary terms. This is beneficial for practitioners because of its ease of use in discussions with top management and reporting with the regulator. In fact, VaR is used by commercial banks, investment banks, pension funds, hedge funds and other financial institutions.

The VaR with confidence level α is given by:

$$VaR(\alpha) = F^{-1}(1-\alpha) \tag{34}$$

where F^{-1} is the inverse of the cumulative distribution function of the returns of a position.

From a probabilistic perspective,

$$Prob(r_t \leq -VaR) = 1 - \alpha \tag{35}$$

where r_t is the return realized by a financial instrument at time t and α represents the confidence level.

From a regulatory point of view, the Basel Committee imposed, as minimum capital requirement for banks using VaR, a level of capital not lower than 3 times (multiplicative factor) the VaR. The value at risk measure to be taken into account is the greater value between the previous day VaR or the average of the VaR measures obtained in the last sixty days. The multiplicative factor can be increased to a factor up to 4 in case the outcome of the model proves to be inadequate in predicting losses. The outcome is evaluated through the so-called backtesting: it is the process of comparing the losses predicted by the VaR model to those actually faced by the financial institution over a defined period of time.

The three most common methods used by financial institutions for VaR estimation are the Variance-Covariance Approach, Historical Simulation, and Monte Carlo Simulation.

The simplest form of the Variance-Covariance method, the *delta-normal* approach, is applicable to portfolios composed by assets whose return is a linear function of the returns of risk factors or underlying assets.

The formula for a single stock position is:

$$VaR = MV \cdot k \cdot \sigma \cdot \sqrt{t} \tag{36}$$

where MV is the current market value of the position, k is the multiplicative factor related to the number of standard deviation according to α, σ is the annual volatility, that is scaled by the square root of time $\sqrt{t}$. For the regulatory purposes of $\alpha = 99\%$, under the normal distribution assumption, k equals 2.33.

In more general terms, if the asset is not a stock, we may add a δ in the formula to account for the sensitivity of the position's market value to changes in the market factor. The formula will thus be: $VaR = MV \cdot \delta \cdot k \cdot \sigma \cdot \sqrt{t}$.

When we shift from a single position to a portfolio, the correlation among the different instruments must be taken into account. $\mathrm{VaR_p}$ of a portfolio is thus:

$$VaR_p = \sqrt{\sum_{i=1}^{N}\sum_{j=1}^{N} VaR_i VaR_j \rho_{i,j}} \tag{37}$$

where $\rho_{i,j}$ is the correlation coefficient between two financial instruments i and j.

In the historical simulation approach, future price changes are assumed to be properly represented by their empirical historical distribution. The rationale is that past behavior reflects the future behavior of market variables. Therefore one will select a proper time frame and cut the distribution of returns at the desired percentile to obtain the VaR estimate.

The last standard approach for Value at Risk measurement is the Monte Carlo simulation. It consists in generating random data given a defined probability density function (the one that best fits the data of returns), and, using the same percentile logic as the historical simulation, rearrange observations and cut them at the desired percentile to find the VaR. A large number of simulations is needed in order to minimize sampling variability due to randomization. As a matter of fact, the Law of Large Numbers ensures convergence to a stable VaR output when the number of iterations increase to a very high number.

When adopting the Variance-Covariance approach and the Monte Carlo simulation approach, an element of crucial importance is the estimation of volatility. In fact, we need volatility as an input both to use the formula that gives us the VaR under the Variance-Covariance approach, as well as to have the parameters needed to generate random scenarios under the Monte Carlo simulation approach.

There are basically three methods to estimate volatility. The first is to use historical volatility, thus assuming that volatility is a constant parameter. The second possibility is to construct a more sophisticated model that allows volatility to change over time. Past volatility here is used to construct the volatility estimation model, but does not coincide with it. Different classes of models exist, from simple moving average (MA) processes to exponentially weighted moving average (EWMA) processes, to generalized autoregressive conditional heteroskedasticity (GARCH) processes. The third method consists in using the volatility implied in option prices. In this case, instead of using a model with volatility as an input to find the option price, the option price is used as an input to find the implied volatility that justifies such a price.

As I have already mentioned, the Basel Committee requires that banks using internal models for market risk backtest the model adopted to check its validity. In particular, they should ensure that the model's predicted losses are consistent with the losses actually realized. The standard tests imposed by the regulator evaluate the accuracy of the model, and there are mainly two types of tests: the Unconditional Coverage Test and the Conditional Coverage Test. Backtesting tests may be used also to compare loss functions among different models. However, the Basel Committee is only interested in the first approach for the purpose determining capital requirements.

The Unconditional Coverage Test was developed by Kupiec in 1995. The null hypothesis to be tested is that the frequency of empirical exceptions is consistent with the confidence level adopted, which reflects the number of "theoretical" exceptions α.

In 1998 Christoffersen developed a Conditional Coverage Test, in order to solve the problem of the Unconditional Coverage Test, namely the fact that it does not take into account the time distribution of exceptions. Christoffersen introduced a test for serial dependence of exceptions.

After discussing the main concepts related to VaR, I presented a brief summary of the vast literature that concerns Value at Risk. I mainly analyzed the methodolgy (data and time period used, VaR models investigated, etc.) and the results of empirical studies that focus on the performance of different VaR models. Although the numerous researches carried out in the past years, no consensus exists among researchers or institutions with respect to the best VaR approach. Among the number of empirical studies carried out by researchers on this topic, I predominantly discussed those that were more relevant for the purpose of the thesis.

In the final section of my work, I described the methodology that I have employed to carry out the empirical study and show the relevant results that I have obtained. Specifically, I divided the chapter into four subsections: in the first I presented the data that I have used to build the portfolio, in the second I reported the sample period taken into account and the framework for the evaluation of the models, in the third I detailed the construction of the above mentioned models, and in the fourth I showed the results of the study.

I built a portfolio composed of five different asset classes, specifically the US stock index NASDAQ Composite (ˆIXIC) that includes almost all the stocks listed on the Nasdaq stock exchange and is heavily weighted towards companies in the IT sector, a fixed income security issued by the government of Italy on 03/02/1997 and with maturity 01/11/2026, paying a fixed semiannual coupon equal to 7.5%, the commodity index PIMCO CommoditiesPLUS Strategy Fund (PCLIX), plus the Bitcoin/USD exchange rate and the EUR/USD exchange rate. I chose this composition in order to have a well-diversified (not only in terms of asset class but also in terms of geopgraphy) and liquid portfolio, with a US stock index and a European fixed income security, plus commodities, FX and a cryptocurrency.

I obtained dollar daily adjusted closing prices (not taking into account splits and dividends) for the five assets for the period from 17/08/2020 to 18/08/2022, thus obtaining 500 observations. It is important to mention that this time period has to be considered as a crisis period, because of the Covid pandemic and the Ukrainian war, that highly increased the volatility of returns, and, for some assets, also deteriorated their financial performance. I constructed a $10 million portfolio by considering the fact that with riskier securities (such as the Bitcoin) having a greater proportion in the portfolio would be too risky, so the portfolio was built in the following way: 25% of the total amount invested in ˆIXIC, another 25% invested in the Italian bond, then 20% in the PCLIX, 10% in the Bitcoin/USD and 20% in the EUR/USD exchange rate.

The 2-year data that I obtained were split so that the first 250 observations served as in-the-sample to generate VaR estimates that were then backtested on the second year of data. For all asset classes, the daily VaR was constructed through a rolling window that included the last 250 observations. In other words, I do not have the Value at Risk for the first 250 observations, which are needed for the estimation of the VaR for the 251th business day and on. VaR was tested at 1% confidence level one-step ahead, so that for a robust VaR model I expected 2 or 3 exceptions over the 250 out-of-the-sample observations. I performed both the Unconditional Coverage Test (Kupiec test) and the Conditional Coverage Test (Christoffersen test) first for each security and then at portfolio aggregate (of course I did not perform the Conditional Coverage test when there were no exceptions at asset level or portfolio level). At aggregate portfolio level, I performed backtesting for two different measures: the portfolio Net VaR, which takes into account the benefit of diversification, and the portfolio Gross VaR, simply obtained by summing the individual VaRs of the five securities. The difference between Gross VaR and Net VaR allows to quantify in monetary terms the diversification effect.

For the purpose of the study, I built VaR models through simulation methods, both historical simulation and Monte Carlo simulation to evaluate the performance of these approaches in the assessment of market risk for a portfolio of securities belonging to five different asset classes. After having computed the Value at Risk estimates, I performed the backtesting to assess the validity of the models. Then, I detailed the procedure I followed to develop the VaR simulations.

For the historical simulation VaR, I cut the distribution of the daily PnL at the first percentile to obtain the daily VaR estimate. For the Monte Carlo simulation VaR, I have assumed that different processes regulate the evolution of the path of the prices of different assets; for example, I assumed that stock prices follow a Geometric Brownian Motion model, while for the fixed income security in the portfolio I applied the Vasicek model.

I find that, if taking into account the securities that compose the portfolio one by one, neither of the two methods are able to prove their validity for all the instruments in such a period of stress. In particular, exceptions are too numerous for the ˆIXIC Index and for the Italian Government bond under both approaches, and for the latter security the outcome is catastrophic with the Monte Carlo simulation. However, when I analyzed the VaR estimate at portfolio aggregate, I witnessed to an improvement in the risk assessment: in fact, the Net VaR computed with historical simulation gave rise to 4 exceptions, which is an acceptable result, in particular in a stress period, and would locate the bank in the green zone under the Basel Accord penalty zones. This result is of important because it grants that if banks are able to diversify their portfolios, they are able to generate appropriate VaR estimates even in a period of stress and even if for some securities inside the portfolio the risk estimation is not optimal. The Gross VaR estimated under historical simulation and Monte Carlo simulation generated, respectively, 1 and 0 exceptions, but in this case one might argue that the level of capital put aside is too high and fails to consider the benefits of diversification into different asset classes. In general, historical simulation performed better than Monte Carlo simulation as for the Unconditional Coverage test. In terms of conditionality of exceptions, neither the individual securities nor the portfolio aggregate are able to pass the Conditional Coverage test and this is dangerous because clustered exceptions may lead to the erosion of the capital of a bank.